arXiv:1604.05757v1 [cs.DM] 19 Apr 2016

Forbidden Subgraph Bounds for Parallel Repetition and the Density Hales-Jewett Theorem

Jan Hązła* Thomas Holenstein† Anup Rao‡

April 19, 2016

Abstract

We study a special kind of bounds (so called forbidden subgraph bounds, cf. Feige, Verbitsky '02) for parallel repetition of multi-prover games.

First, we show that forbidden subgraph upper bounds for $r \geq 3$ provers imply the same bounds for the density Hales-Jewett theorem for alphabet of size r. As a consequence, this yields a new family of games with slow decrease in the parallel repetition value.

Second, we introduce a new technique for proving exponential forbidden subgraph upper bounds and explore its power and limitations. In particular, we obtain exponential upper bounds for two-prover games with question graphs of treewidth at most two and show that our method cannot give exponential bounds for all two-prover graphs.

1 Introduction

1.1 Multi-prover games

An *r-prover game* is a protocol in which r *provers* have a joint objective of making another entity, the *verifier*, accept. The execution of such a game looks as follows: The verifier first samples r questions $q^{(1)}, \ldots, q^{(r)} \in Q^{(1)} \times \ldots \times Q^{(r)}$. Those questions are sampled uniformly from some *question set* $\overline{Q} \subseteq Q^{(1)} \times \ldots \times Q^{(r)}$.

Then, she sends the questions to the provers: the j-th prover receives $q^{(j)}$ and sends back an answer $a^{(j)}$ (from a finite answer alphabet $A^{(j)}$) that depends only on $q^{(j)}$. Finally, the verifier accepts or rejects based on the evaluation of a *verification predicate* $V(q^{(1)}, \ldots, q^{(r)}, a^{(1)}, \ldots, a^{(r)})$.

*ETH Zürich, Department of Computer Science, Zurich, Switzerland. E-mail: jan.hazla@inf.ethz.ch. J. H. is supported by the Swiss National Science Foundation (SNF), project no. 200021-132508.

†Google, Zurich, Switzerland. E-mail: thomas.holenstein@gmail.com. Part of this work was done while T. H and A. R were at the Simons Institute.

‡University of Washington, Seattle, USA. E-mail: anuprao@cs.washington.edu

A *strategy* $(\mathcal{S}^{(1)}, \ldots, \mathcal{S}^{(r)})$ for the provers consists of r functions, the j-th of which maps questions to answers for the j-th prover. The *value* of a game is

$$\mathrm{val}(\mathcal{G}) := \max_{\mathcal{S}^{(1)},\ldots,\mathcal{S}^{(r)}} \Pr\left[V\left(q^{(1)}, \ldots, q^{(r)}, \mathcal{S}^{(1)}(q^{(1)}), \ldots, \mathcal{S}^{(r)}(q^{(r)})\right) = 1\right] ,$$

where maximum is over all strategies and the probability over the uniform choice of $q^{(1)}, \ldots, q^{(r)} \in \overline{Q}$. A game $\mathcal{G}$ is called *trivial* if $\mathrm{val}(\mathcal{G}) = 1$.

A formal definition is provided in Section 2.2. One might consider allowing other distributions over $\overline{Q}$ than the uniform one. This does not make much difference for parallel repetition, as shown in Section 2.3.

1.2 Parallel repetition

The *n-fold parallel repetition* $\mathcal{G}^n$ of an r-prover game $\mathcal{G}$ is a game where the verifier samples n independent question tuples, sends n questions to each prover, receives n answers from each prover, and accepts if all n instances of the verification predicate for $\mathcal{G}$ accept.

It is easy to see that $\mathrm{val}(\mathcal{G}^n) \geq \mathrm{val}(\mathcal{G})^n$, in particular if $\mathcal{G}$ is trivial, then $\mathcal{G}^n$ is trivial as well. However, since the i-th answer of a prover can depend on all of his questions, as opposed to just the i-th one, it is possible that $\mathrm{val}(\mathcal{G}^n)$ attains a higher value.

In this paper we are interested in upper bounds on $\mathrm{val}(\mathcal{G}^n)$ that depend only on $\overline{Q}$ and n. They are called *forbidden subgraph* bounds, with the name explained in [FV02].

Specifically, let $\omega_{\overline{Q}}(n) := \max_{\mathcal{G}} \mathrm{val}(\mathcal{G}^n)$, where the maximum is over all non-trivial games $\mathcal{G}$ with question set[1] $\overline{Q}$. We say that a question set $\overline{Q}$ *admits parallel repetition* if $\lim_{n\to\infty} \omega_{\overline{Q}}(n) = 0$. Furthermore, we will say that $\overline{Q}$ *admits exponential parallel repetition* if there exists $C_{\overline{Q}} < 1$ such that

$$\omega_{\overline{Q}}(n) \leq (C_{\overline{Q}})^n .$$

Background The two-prover games in the context of theoretical computer science were first introduced by Ben-Or, Goldwasser, Kilian and Wigderson [BGKW88]. Fortnow, Rompel and Sipser were the first to treat the value of two-prover repeated games and exhibit an example with $\mathrm{val}(\mathcal{G}^2) > \mathrm{val}(\mathcal{G})^2$ [FRS88, FRS90]. Extensive works on parallel repetition were produced, for a survey we refer to [Fei95] and [Raz10]. Here we will only mention some results relevant to our theorems.

1.3 Density Hales-Jewett theorem

Definition 1.1 (Combinatorial line)**.** Let $r, n \in \mathbb{N}_{>0}$ and $[r] := \{1, \ldots, r\}$. A *combinatorial pattern* over $[r]^n$ is a string

$$(b_1, \ldots, b_n) = \underline{b} \in ([r] \cup \{\star\})^n \setminus [r]^n ,$$

[1] Note that the number of provers r is implicitly determined by $\overline{Q}$.

where $\star$ is a special symbol called the *wildcard*. Note that a pattern contains at least one wildcard.

For $q \in [r]$ we let $\underline{b}(q) \in [r]^n$ to be the string formed from $\underline{b}$ by substituting all occurrences of the wildcard with q.

A *combinatorial line* associated with a pattern $\underline{b}$ is the set $L(\underline{b}) := \{\underline{b}(1), \ldots, \underline{b}(r)\}$. $\Diamond$

Example 1.2. For $r = 3$, $n = 5$, an example pattern is $\underline{b} = 12\star2\star$. The corresponding combinatorial line is $L(\underline{b}) = \{12121, 12222, 12323\}$. $\Diamond$

The Hales-Jewett theorem [HJ63] says that for any r, and large enough n, any coloring of $[r]^n$ contains a monochromatic combinatorial line. The density Hales-Jewett theorem is a strengthening of this result: It states that for any r and μ, and large enough n, any subset of $[r]^n$ of measure μ contains a combinatorial line.

These theorems have played a central role in Ramsey theory because many other results in this field (e.g., corners theorem [AS74], Szemerédi's theorem [Sze75]) can be reduced to the density Hales-Jewett theorem.

Definition 1.3. For a set $S \subseteq [r]^n$ we define its *measure* as $\mu(S) := |S|/r^n$. We then let $\omega_r^{\mathrm{DHJ}}(n)$ to be the maximum measure of a subset of $[r]^n$ that does not contain a combinatorial line. $\Diamond$

Theorem 1.4 (Density Hales-Jewett theorem, [FK91])**.** *Let $r \geq 2$. Then,*

$$\lim_{n\to\infty} \omega_r^{\mathrm{DHJ}}(n) = 0 \,.$$

The original proof of Furstenberg and Katznelson [FK91] does not give explicit bounds for $\omega_r^{\mathrm{DHJ}}(n)$. This situation was improved by a more recent proof by Polymath [Pol12], however their bounds are still not primitive recursive: They prove that $\omega_r^{\mathrm{DHJ}}(n)$ decreases at a rate that is related to the inverse of the r-th Ackermann function of n. In particular, $\omega_3^{\mathrm{DHJ}}(n) \leq O\left(1/\sqrt{\log^* n}\right)$.

On the other hand, another paper by Polymath [Pol10] gives the best known density Hales-Jewett lower bounds: $\omega_r^{\mathrm{DHJ}}(n) \geq \exp\left(-O(\log n)^{1/\lceil \log_2 r \rceil}\right)$ for $r \geq 3$. Note that $r = 2$ is a somewhat special case with $\omega_2^{\mathrm{DHJ}}(n) = \Theta(1/\sqrt{n})$ known by Sperner's theorem.

Verbitsky [Ver96] used the density Hales-Jewett theorem to show that every question set admits parallel repetition:

Theorem 1.5 ([Ver96])**.** *Let $\overline{Q}$ be a k-prover question set of size $\left|\overline{Q}\right| = r$. Then,*

$$\omega_{\overline{Q}}(n) \leq \omega_r^{\mathrm{DHJ}}(n) \,.$$

In particular, $\overline{Q}$ admits parallel repetition.

1.4 Our results — equivalence of DHJ and PR

Despite numerous works, especially concerning the two-prover case, Theorem 1.5 remains the best available bound for general parallel repetition of multi-prover games. Our main result hints that there might be a reason for this situation.

Definition 1.6. Let $r \geq 2$. We define an r-prover question set $\overline{Q}_r \subseteq \{0,1\}^r$ of size r, where the j-th question contains 1 in the j-th position and 0 in the remaining positions. In other words,

$$\overline{Q}_r := \left\{ (q^{(1)}, \ldots, q^{(r)}) : |\{j : q^{(j)} = 1\}| = 1 \right\} .$$

$\Diamond$

Our main theorem is a construction that ties the existence of a combinatorial line in a set with the parallel repetition value of a certain game. Note that it requires at least three provers.

Theorem 1.7. *Let $r \geq 3$, $n \geq 1$ and $S \subseteq [r]^n$ with $\mu(S) = |S|/r^n$ such that S does not contain a combinatorial line.*

There exists an r-prover game $\mathcal{G}_S$ with question set $\overline{Q}_r$ and with answer alphabets $A^{(j)} = 2^{[n]} \times [n]$ such that:

- $\mathrm{val}(\mathcal{G}_S) \leq 1 - 1/r$.
- $\mathrm{val}(\mathcal{G}_S^n) \geq \mu(S)$.

Theorem 1.7 immediately implies an inequality complementary to Theorem 1.5:

Theorem 1.8. *Let $r \geq 3$. We have $\omega_r^{\mathrm{DHJ}}(n) \leq \omega_{\overline{Q}_r}(n)$.*

Our game construction is related to another one by Feige and Verbitsky [FV02] in the following way: They proved that the so-called forbidden subgraph method is universal[2] for obtaining bounds on $\omega_{\overline{Q}}(n)$ (hence the name: forbidden subgraph bounds). In their proof they used a result that is very similar to Theorem 1.7 generalized to an arbitrary question set $\overline{Q}$: its statement is almost the same except they assume that the set S does not contain a forbidden subgraph (instead of a combinatorial line). As a matter of fact, our game $\mathcal{G}_S$ is the same as the one from their theorem instantiated on the question set $\overline{Q}_r$.

Consequently, Theorem 1.7 has a shorter proof assuming the result of Feige and Verbitsky: One only checks that for the question set $\overline{Q}_r$ a forbidden subgraph in S implies a combinatorial line in S. This is discussed in more detail in Section 3.1. A self-contained proof of Theorem 1.7 is provided in Section 3.

[2] As a matter of fact, they treated only two-prover games. However, their construction generalized to the multi-prover setting.

Most of the work done on parallel repetition lower bounds is based on two-prover examples by Feige and Verbitsky [FV02] and Raz [Raz11] with little known in the multi-prover case. By taking the largest known combinatorial line-free sets from [Pol10], Theorem 1.7 yields a new family of multi-prover games with a slow decrease in the parallel repetition value. In particular, it establishes the first known question sets that do not admit exponential parallel repetetition answering an open question asked in [FV02] in multi-prover case:

Theorem 1.9. *Let $r \geq 3$. The question set $\overline{Q}_r$ does not admit exponential parallel repetition.*

Furthermore, this family is related to another aspect of the lower bound in [FV02]. They show that for any two-prover upper bound of the form

$$\mathrm{val}(\mathcal{G}^n) \leq \exp\left(-f(\epsilon) \cdot g(\left|\overline{A}\right|) \cdot n\right) \ ,$$

where $\mathrm{val}(\mathcal{G}) = 1 - \epsilon$ and $\left|\overline{A}\right|$ is the answer set size, it must be $g(\left|\overline{A}\right|) \leq O\big(\log\log\left|\overline{A}\right| / \log\left|\overline{A}\right|\big)$. On the other hand, the famous two-prover upper bound by Raz [Raz98] has $g(\left|\overline{A}\right|) \geq \Omega\left(1/\log\left|\overline{A}\right|\right)$. The analysis of our family of games implies that if such bound is ever extended to the multi-prover setting, it must be $g\left(\left|\overline{A}\right|\right) \leq O\left(\left(\log\log\left|\overline{A}\right|\right)^{\epsilon} / \log\left|\overline{A}\right|\right)$ for any $\epsilon > 0$.

Those lower bounds are discussed in Section 4.

Finally, inspired by the proof of Theorem 1.7, in Section 5 we define *coloring games* and their parallel repetition and show that it is equivalent to the (coloring) Hales-Jewett theorem.

1.5 Our results — upper bounds on parallel repetition

It remains open if all two-prover question sets admit exponential parallel repetition. We believe that this question, though somewhat forgotten, becomes more interesting in light of Theorem 1.9. If the answer is affirmative, it would constitute a significant difference in behavior of two and three-prover games. On the other hand, if there exist "hard" two-prover sets, it might be possible to use them to obtain non-trivial consequences in vein of Theorem 1.8.

Motivated by this, we explore a new method for obtaining exponential parallel repetition: In Section 6.1 we define a notion of a question set $\overline{Q}$ that is *constructible by conditioning* (see Definition 6.2). It is an inductive definition using two graph-theoretic operations (we interpret an r-prover question set $\overline{Q}$ as the edge set of an r-regular, r-partite hypergraph). Then, using a technique inspired by a previous paper of some of the authors [HHM15], in Section 6.2 we prove:

Theorem 1.10. *Let $\overline{Q}$ be an r-prover question set that is constructible by conditioning. Then, $\overline{Q}$ admits exponential parallel repetition.*

In particular, this gives a new proof of exponential parallel repetition for *free* multi-prover question sets[3] [CCL92, Fei95, Pel95], i.e., those, where the provers' questions are independent.

In Section 7 we present an example of what can be achieved with this method. We show:

Theorem 1.11. *Every bipartite graph G with treewidth at most two is constructible by conditioning. In particular, if G is interpreted as a two-prover question set, then it admits exponential parallel repetition.*

This improves on previous work by Verbitsky [Ver95] and Weissenberger[4] [Wei13] which showed exponential parallel repetition for, respectively, trees and cycles. We note that our proof dooes not seem to be a generalisation of the earlier ones, but rather a genuinely new approach.

In Section 8 we show that our technique of graph constructability is too weak to resolve the two-prover question positively:

Theorem 1.12. *There exists a bipartite graph that is not constructible by conditioning.*

1.6 Comparison with information-theoretic bounds

The general two-prover bound by Raz [Raz98] as improved by Holenstein [Hol09] gives

$$\mathrm{val}(\mathcal{G}^n) \leq \exp\left(-\Omega\left(\epsilon^3 n / \log\left|\overline{A}\right|\right)\right) \tag{1}$$

for a game $\mathcal{G}$ with $\mathrm{val}(\mathcal{G}) = 1 - \epsilon$ and answer set $\overline{A} = A^{(1)} \times A^{(2)}$. There are numerous other bounds of this form (i.e., with dependence in the exponent on some power of ϵ and possibly $\left|\overline{A}\right|$), in particular for projection and free two-prover games [Rao11, BRR$^+$09] and for multi-prover games with quantum entanglement or no-signaling strategies, e.g., [BFS14, CWY15].

On the other hand, we are interested in forbidden subgraph bounds that depend only on the question set $\overline{Q}$. Those two types of bounds are not comparable to each other. Furthermore, it seems that the respective proof techniques are also quite different: while for the bounds like (1) information theory is usually employed, the forbidden graph bounds use more combinatorial arguments.

We do not know if an "information-theoretic" bound holds for general games with more than two provers (it is known only for free games, see [CWY15], and for the so called *anchored games* [BVY15]). We stress that our negative result in Theorem 1.9 does *not* exclude the possibility of such a bound.

[3] The proof in [CCL92] is for two provers only, but it generalizes to many provers by using a result on hypergraphs by Erdős [Erd64].

[4] As a matter of fact, the result of Weissenberger is even stronger: it gives an exponential upper bound that depends only on $\mathrm{val}(\mathcal{G})$, but not on question or answer set size.

2 Preliminaries

2.1 Notation

Many of our results feature two-dimensional vectors. We adopt the following conventions:

- Most of the time we consider two dimensions corresponding to n independent coordinates and r provers. Usually n is meant to be large compared to r.
- We index the n-dimension with i in the subscript and the r-dimension with j in parentheses in the superscript. We denote aggregation over i by underline and over j by overline. For example:

$$\overline{\underline{V}} = (\underline{V}^{(1)}, \ldots, \underline{V}^{(j)}, \ldots, \underline{V}^{(r)}) = (\overline{V}_1, \ldots, \overline{V}_i, \ldots, \overline{V}_n),$$
$$\underline{V}^{(j)} = (V_1^{(j)}, \ldots, V_i^{(j)}, \ldots, V_n^{(j)}), \quad \overline{V}_i = (V_i^{(1)}, \ldots, V_i^{(j)}, \ldots, V_i^{(r)})$$

- We call the element collections aggregated over i (like $\underline{v}^{(j)}$) *vectors* and the element collections aggregated over j (like $\overline{v}_i$) *tuples*.

For sets A, B we sometimes denote $A \cup B$ as $A \dot{\cup} B$ to emphasize that $A \cap B = \emptyset$. For an event $\mathcal{E}$ we denote its indicator function by $\mathbb{1}_{\mathcal{E}}$. Whenever we speak of a partition of a set, we allow empty classes in the partition. The powerset of X is denoted by 2^X. In accordance with the computer science tradition, the symbol log denotes the logarithm with base two. For a string $\underline{x} \in X^n$ and $y \in X$ we let $w_y(\underline{x}) := |\{i \in [n] : x_i = y\}|$.

2.2 Definitions

In this section we provide formal definitions of the most important concepts we use.

Definition 2.1 (Multi-prover games)**.** An *r-prover game* $\mathcal{G} = (\overline{Q}, \overline{A}, V)$ consists of the following elements:

- $\overline{Q} \subseteq Q^{(1)} \times \ldots \times Q^{(r)}$ is a finite *question set*.

  Note that $\overline{Q}$ does not have to consist of all possible tuples. However, we will always assume that there are no "impossible questions", i.e., that for each element of a *question alphabet* $q^{(j)} \in Q^{(j)}$ there exists at least one question tuple $\overline{q} \in \overline{Q}$ with $q^{(j)}$ as its j-th element.
- $\overline{A} = A^{(1)} \times \ldots \times A^{(r)}$ is a finite *answer set*.
- $V : \overline{Q} \times \overline{A} \to \{0, 1\}$ is a *verification predicate*.

A *strategy* $\overline{\mathcal{S}} = (\mathcal{S}^{(1)}, \ldots, \mathcal{S}^{(r)})$ for a game $\mathcal{G}$ is a tuple of functions, where $\mathcal{S}^{(j)} : Q^{(j)} \to A^{(j)}$.

Let $\overline{q} = (q^{(1)}, \ldots, q^{(r)})$ be a random variable sampled uniformly from $\overline{Q}$. For a strategy $\overline{\mathcal{S}}$ let $\overline{\mathcal{S}}(\overline{q}) := (\mathcal{S}^{(1)}(q^{(1)}), \ldots, \mathcal{S}^{(r)}(q^{(r)}))$.

We define the *value* of a game $\mathcal{G}$ as

$$\mathrm{val}(\mathcal{G}) := \max_{\overline{\mathcal{S}}} \Pr\left[V\left(\overline{q}, \overline{\mathcal{S}}(\overline{q})\right) = 1\right] .$$

We say that a game is *trivial* if its value is 1.

We also say it is *free* if $\overline{Q} = Q^{(1)} \times \ldots \times Q^{(r)}$. Note that in our setting this is equivalent to the property that the provers' questions are distributed independently. ◊

Definition 2.2 (Parallel repetition)**.** The *n-fold parallel repetition* $\mathcal{G}^n$ of an r-prover game $\mathcal{G} = (\overline{Q}, \overline{A}, V)$ is another r-prover game $\mathcal{G}^n = \left(\underline{\overline{Q}}, \underline{\overline{A}}, \underline{V}\right)$ where

- The question alphabet for the j-th prover $\underline{Q}^{(j)} := \left(Q^{(j)}\right)^n$ is the n-fold product of the original $Q^{(j)}$. Consequently, the question set $\underline{\overline{Q}}$ is the n-fold product of $\overline{Q}$.
- In the same way, the answer alphabet for the j-th prover $\underline{A}^{(j)}$ is the n-fold product of $A^{(j)}$.
- The verification predicate $\underline{V}$ accepts if and only if all of its n single instances accept:

$$\underline{V}\left(\underline{\overline{q}}, \underline{\overline{a}}\right) = 1 \iff \forall i \in [n] : V\left(\overline{q}_i, \overline{a}_i\right) = 1 .$$

◊

Definition 2.3. For an r-prover question set $\overline{Q}$ we define $\omega_{\overline{Q}}(n) := \max_{\mathcal{G}} \mathrm{val}(\mathcal{G}^n)$, where the maximum is over all non-trivial games $\mathcal{G}$ with question set $\overline{Q}$.

We say that $\overline{Q}$ *admits parallel repetition* if $\lim_{n\to\infty} \omega_{\overline{Q}}(n) = 0$. We say that $\overline{Q}$ *admits exponential parallel repetition* if there exists $C_{\overline{Q}} < 1$ such that for every $n \in \mathbb{N}$:

$$\omega_{\overline{Q}}(n) \leq (C_{\overline{Q}})^n .$$

◊

An important notion in our proofs is a homomorphism of r-regular, r-partite hypergraphs:

Definition 2.4. Let $r \geq 2$ and $\overline{Q} \subseteq Q^{(1)} \times \ldots \times Q^{(r)}$ be an r-prover question set. Consider the r-regular, r-partite hypergraph $G = (Q^{(1)}, \ldots, Q^{(r)}, \overline{Q})$. We will often abuse the notation by identifying this hypergraph with $\overline{Q}$.

Given two hypergraphs $(Q^{(1)}, \ldots, Q^{(r)}, \overline{Q})$ and $(P^{(1)}, \ldots, P^{(r)}, \overline{P})$ we say that $f = (f^{(1)}, \ldots, f^{(r)})$, $f^{(j)} : Q^{(j)} \to P^{(j)}$ is a *homomorphism* from $\overline{Q}$ to $\overline{P}$ if $\overline{q} = (q^{(1)}, \ldots, q^{(r)}) \in \overline{Q}$ implies $f(\overline{q}) := (f^{(1)}(q^{(1)}), \ldots, f^{(r)}(q^{(r)})) \in \overline{P}$.

If f is a homomorphism from $\overline{Q}$ to $\overline{Q}$ we will just say that f is a homomorphism of $\overline{Q}$. We denote the set of homomorphisms from $\overline{Q}$ to $\overline{P}$ by $\mathrm{Hom}(\overline{Q}, \overline{P})$. ◊

There are some homomorphisms of an r-prover question set $\overline{Q}$ that are important to us:

Definition 2.5. We denote the *identity homomorphism* as Id. We will also use the *constant homomorphism* $\mathbb{1}_{\overline{q}}$ for a hyperedge $\overline{q} \in \overline{Q}$, where $\mathbb{1}_{\overline{q}}$ maps every hyperedge to $\overline{q}$. $\Diamond$

2.3 Reduction of general parallel repetition to uniform case

We show how parallel repetition for a question distribution that is not necessarily uniform over a question set $\overline{Q}$ reduces to the uniform case. The proof is taken from [FV02] and is included here for completeness.

Theorem 2.6. *Let $\overline{Q}$ be an r-prover question set and let $\mathcal{Q}$ be a probability distribution with support $\overline{Q}$ such that $\epsilon := \min_{\overline{q} \in \overline{Q}} \mathcal{Q}(\overline{q})$ and*

$$\alpha := \frac{\epsilon}{1/\left|\overline{Q}\right|} = \epsilon \left|\overline{Q}\right| \ .$$

Furthermore, assume that a function $f : \mathbb{N}_{>0} \to [0,1]$ is such that for every non-trivial game $\mathcal{H}$ uniform over $\overline{Q}$:

$$\mathrm{val}(\mathcal{H}^n) \leq f(n) \ .$$

Then, for every non-trivial game $\mathcal{G}$ such that its questions are sampled according to $\mathcal{Q}$ we have

$$\mathrm{val}(\mathcal{G}^n) \leq \exp\left(-\alpha^2 n/2\right) + f\left(\alpha n/2\right) \ .$$

Proof. First, note that we can write $\mathcal{Q} = \alpha U_{\overline{Q}} + (1-\alpha)\mathcal{Q}'$, where $U_{\overline{Q}}$ is the uniform distribution over $\overline{Q}$ and $\mathcal{Q}'$ some other probability distribution. Consequently, we can define an i.i.d. random binary vector $\underline{B} = (B_1, \ldots, B_n)$ coupled with an execution of $\mathcal{G}^n$ such that $B_i = 1$ if the i-th question is sampled from $U_{\overline{Q}}$ and $B_i = 0$ if it was sampled from $\mathcal{Q}'$.

Consider an execution of $\mathcal{G}^n$ with a modified verifier. The new verifier first checks if $w_1(\underline{B}) \geq \alpha n/2$, i.e., if the number of coordinates with $B_i = 1$ is at least half of the expectation αn. She accepts if this check fails. If the first check succeeds, the new verifier accepts if the single-coordinate verifier accepts on all coordinates with $B_i = 1$.

Let us call this modified game $(\mathcal{G}^n)'$. Clearly, $\mathrm{val}(\mathcal{G}^n) \leq \mathrm{val}((\mathcal{G}^n)')$. Furthermore, let $\mathcal{G}^*$ be a game with the same verifier as $\mathcal{G}$ but uniform over $\overline{Q}$. Note that $\mathcal{G}^*$ is non-trivial. Observe that conditioned on a choice of $\underline{B}$, the game $(\mathcal{G}^n)'$ is the same as $\mathcal{G}^*$ repeated $w_1(\underline{B})$ times. Consequently, and using Chernoff bound,

$$\begin{aligned}\mathrm{val}(\mathcal{G}^n) \leq \mathrm{val}((\mathcal{G}^n)') &= \mathrm{E}\left[\mathrm{val}((\mathcal{G}^n)') \mid \underline{B}\right] \leq \Pr\left[w_1(\underline{B}) < \alpha n/2\right] + f(\alpha n/2) \\ &\leq \exp(-\alpha^2 n/2) + f(\alpha n/2) \ .\end{aligned}$$

□

3 Parallel Repetition Implies Density Hales-Jewett

In this section we prove Theorem 1.7. We present a self-contained proof here and then we explain how it is related to the proof from [FV02] in Section 3.1.

Let $r \geq 3$, $n \geq 1$ and $S \subseteq [r]^n$ with $\mu(S) = |S|/r^n$. We want to define a game $\mathcal{G}_S$ with question set $\overline{Q}_r$ such that:

- If S does not contain a combinatorial line, then $\mathcal{G}_S$ is non-trivial.
- $\mathrm{val}(\mathcal{G}_S^n) \geq \mu(S)$.

Firstly, note that there is a natural bijection between the question tuples in $\overline{Q}_r$ and $[r]$. Namely, we can think of the verifier as choosing the number of a *special* prover $a \in [r]$ u.a.r. and sending 1 to the special prover and 0 to all other provers.

The answer alphabet of the game $\mathcal{G}_S$ is the same for all provers: $A^{(j)} := 2^{[n]} \times [n]$. Upon sampling a special prover a and receiving answers $(T^{(1)}, z^{(1)}), \ldots, (T^{(r)}, z^{(r)})$, the verifier checks the following conditions and accepts if all of them are met:

- The sets $T^{(1)}, \ldots, T^{(r)}$ form a partition of $[n]$.
- $z^{(1)} = z^{(2)} = \ldots = z^{(r)} = z$.
- $z \in T^{(a)}$.
- Let $\underline{s} = (s_1, \ldots, s_n)$ be the string over $[r]^n$ such that $s_i = j$ iff $i \in T^{(j)}$. Then, $\underline{s} \in S$.

The next claim is not actually needed for the proof. However, it provides some intuition for the construction of the game.

Claim 3.1. *If S has a combinatorial line, then the game $\mathcal{G}_S$ is trivial.*

Proof. Let $\underline{b}$ be a pattern with its combinatorial line $L(\underline{b}) \subseteq S$, and fix a position $z \in [n]$ with $b_z = \star$. For $\sigma \in [r] \cup \{\star\}$, let $B(\sigma) := \{i : b_i = \sigma\}$, the set of coordinates in which $\underline{b}$ equals σ. We consider the strategy in which prover j responds with

$$\mathcal{P}^{(j)}(q^{(j)}) := \begin{cases} \big(B(j), z\big) & \text{if } q^{(j)} = 0, \\ \big(B(j) \cup B(\star), z\big) & \text{if } q^{(j)} = 1. \end{cases} \tag{2}$$

Since the sets $B(1), \ldots, B(r), B(\star)$ form a partition of $[n]$, the verifier will always accept the first condition. The condition $z^{(1)} = \cdots = z^{(r)}$ is obviously always met. Also $z \in T^{(a)}$ is clear, since prover a responds with $(B(a) \cup B(\star))$ and $z \in B(\star)$. Finally, $\underline{s} \in S$ is also clear since $\underline{s}$ is exactly the pattern $\underline{b}$ with a in place of stars, i.e., $\underline{s} = \underline{b}(a)$ and since $\underline{b}(a) \in L(\underline{b}) \subseteq S$. □

Claim 3.2. *If the game $\mathcal{G}_S$ is trivial, then S has a combinatorial line.*

Proof. Let $\mathcal{P}^{(1)}, \ldots, \mathcal{P}^{(r)}$ be a strategy for the provers that always wins.

For $q \in \{0,1\}$ and $j \in [r]$, we let $(T^{(j)}_{(q)}, z^{(j)}_{(q)}) =: \mathcal{P}^{(j)}(q)$ be the answer which prover j gives on question q. Since the verifier checks $z^{(j)}_{(0)} = z^{(a)}_{(1)}$ whenever $j \neq a$, we see that $z^{(1)}_{(0)} = z^{(2)}_{(0)} = \ldots = z^{(r)}_{(0)} = z^{(1)}_{(1)} = \ldots = z^{(r)}_{(1)} =: z$ (note that we used $r \geq 3$).

Next, for any two $j \neq j'$, the sets $T^{(j)}_{(0)}$ and $T^{(j')}_{(0)}$ are pairwise disjoint. Otherwise, if the verifier chooses a which is different from both j and j', she will reject.

Furthermore, $z \notin T^{(1)}_{(0)} \cup \ldots \cup T^{(r)}_{(0)}$, since if $z \in T^{(j)}_{(0)}$, the verifier rejects if $a \neq j$. Hence, the following defines a combinatorial pattern $\underline{b}$:

$$b_i := \begin{cases} j & \text{if } i \in T^{(j)}_{(0)}, \text{ for } j \in [r], \\ \star & \text{otherwise.} \end{cases} \tag{3}$$

Fix now $a \in [r]$. We show that $\underline{b}(a) \in S$. Suppose the verifier picks a as the special prover. Since the verifier checks that the sets $T^{(\cdot)}$ form a partition, it must be that prover a responds with $T^{(a)}_{(1)} = [n] \setminus \left(T^{(1)}_{(0)} \cup \ldots \cup T^{(a-1)}_{(0)} \cup T^{(a+1)}_{(0)} \cup \ldots T^{(r)}_{(0)}\right)$. Since the verifier checks that the resulting string is in S and accepts, it must be that $\underline{b}(a) \in S$. This holds for every a, and thus $L(\underline{b}) \subseteq S$. □

Claim 3.3. *The value of $\mathcal{G}^n_S$ is at least $\mu(S)$.*

Proof. Let $T^{(j)}$ be the set of coordinates in which prover j is special, $T^{(j)} := \{i \in [n] : q^{(j)}_i = 1\}$. In coordinate i, prover j responds with $(T^{(j)}, i)$.

Let $a_1, \ldots, a_n$ be the sequence of special provers which the verifier picks. We claim that if $(a_1, \ldots, a_n) \in S$ then the verifier accepts in all coordinates. Of course this happens with probability $\mu(S)$.

To see this, note first that the sets $T^{(1)}, \ldots, T^{(r)}$ indeed form a partition of $[n]$ (because in each coordinate there is exactly one special prover). Next, $z^{(1)}_i = \ldots = z^{(r)}_i = i \in T^{(a_i)}$, by definition of $T^{(j)}$ and since prover a_i is special in coordinate i.

Finally, $\underline{s} \in S$, since for all n coordinates $\underline{s}$ is exactly the string $(a_1, \ldots, a_n)$. □

3.1 Connection to the universality proof in [FV02]

We explain the connection between our proof of Theorem 1.7 and [FV02]. Recall our definition of a hypergraph homomorphism (Definitions 2.4 and 2.5).

Definition 3.4. Let $\overline{Q}$ be an r-prover question set and let $\underline{Q} := \overline{Q}^n$ be its n-fold parallel repetition. Let $S \subseteq \underline{Q}$ with $\mu(S) = |S| / \left|\overline{Q}\right|^n$ and let $\underline{f} = (f_1, \ldots, f_n)$ be a vector of n homomorphisms of $\overline{Q}$. We say that $\underline{f}$ is *good* for S if:

- For every $\overline{q} \in \overline{Q}$ we have that $\underline{f}(\overline{q}) := (f_1(\overline{q}), \ldots, f_n(\overline{q})) \in S$.

- There exists $i \in [n]$ such that f_i is identity.

We say that the question set $\overline{Q}$ is (n, ϵ)-good if for every $S \subseteq \underline{\overline{Q}}$ with $\mu(S) \geq \epsilon$ there exists a vector of homomorphisms that is good for S. $\Diamond$

It turns out that if $\overline{Q}$ is (n, ϵ)-good, then $\omega_{\overline{Q}}(n) \leq \epsilon$. This is the forbidden subgraph method and it is presented in Section 6.2.1. Verbitsky [Ver95, FV02] discovered a related game construction:

Theorem 3.5 ([FV02])**.** *Let $\overline{Q}$ be a connected, r-prover question set and $S \subseteq \underline{\overline{Q}}$. There exists an r-prover game $\mathcal{G}_S$ with question set $\overline{Q}$ such that:*

- *If $\mathcal{G}_S$ is trivial, then there exists a homomorphism vector $\underline{f}$ that is good for S.*
- $\mathrm{val}(\mathcal{G}_S^n) \geq \mu(S)$.

Note that Theorem 3.5 implies that the forbidden subgraph method is universal in the sense that it gives the best possible bounds on $\omega_{\overline{Q}}(n)$:

Corollary 3.6 ([FV02])**.** *Let $\overline{Q}$ be a connected, r-prover question set. Then,*

$$\omega_{\overline{Q}}(n) = \inf\left\{\epsilon : \overline{Q} \text{ is } (n,\epsilon)\text{-good}\right\} .$$

There are a couple of caveats with regards to our formulation of Theorem 3.5: First, we state it in terms of homomorphisms instead of forbidden subgraphs as in [FV02]. However, (with hindsight) both statements are easily seen to be equivalent. Second, the proof in [FV02] is only for the two-prover case. However, it generalizes to multiple provers in a natural way.

It turns out that our construction of the game $\mathcal{G}_S$ for Theorem 1.7 is an instantiation of the construction from Theorem 3.5 for the question set $\overline{Q}_r$. Consequently, Theorem 3.5 has a shorter proof that assumes Theorem 1.7. We find it instructive to present it below:

Claim 3.7. *Let $r \geq 3$. The only homomorphisms of the question set $\overline{Q}_r$ are identity and constants.*

Proof. As previously, we identify the edges in $\overline{Q}_r$ with numbers $a \in [r]$. Assume that there exists a homomorphism $f = (f^{(1)}, \ldots, f^{(r)})$ that maps edge b to some $a \neq b$. We show that it must be $f = \mathbb{1}_a$.

First, since b is mapped to a, we must have $f^{(b)}(1) = 0$, $f^{(a)}(0) = 1$ and $f^{(j)}(0) = 0$ for every $j \notin \{a, b\}$. $f^{(a)}(0) = 1$ implies that every edge $j \neq a$ is also mapped to a and, consequently, $f^{(j)}(0) = f^{(j)}(1) = 0$ for every $j \neq a$. But from this it follows that also edge a must be mapped onto itself, hence $f = \mathbb{1}_a$. $\square$

Claim 3.8. *Let $r \geq 3$ and $S \subseteq \underline{\overline{Q}}_r \cong [r]^n$ such that there exists a homomorphism vector $\underline{f}$ that is good for S. Then, S contains a combinatorial line.*

Proof. Define a pattern $\underline{b}$ as

$$b_i := \begin{cases} j & \text{if } f_i = \mathbb{1}_j, \\ \star & \text{if } f_i = \mathrm{Id}. \end{cases}$$

By Claim 3.7, $\underline{b}$ is well-defined. From the second point in a definition of a good vector, there is a star on at least one coordinate. From the first point in that definition, $\underline{b}(j) \in S$ for every $j \in [r]$. Consequently, $L(\underline{b}) \subseteq S$. □

Proof of Theorem 1.7. Let $r \geq 3$ and $S \subseteq [r]^n \cong \overline{Q}_r$ such that S does not have a combinatorial line. By Claim 3.8, there is no homomorphism vector good for S (viewed as a subset of $\overline{Q}_r$). But now the game from Theorem 3.5 does the job. □

4 Lower Bounds on Multi-Prover Parallel Repetition

In this section we explore some lower bounds on parallel repetition implied by Theorem 1.7. Our main observation is Theorem 1.9: for more than two provers there exist question sets that do not admit exponential parallel repetition.

Proof of Theorem 1.9. Let n be divisible by r and let S contain all strings with equidistributed alphabet elements, i.e.,

$$S := \{\underline{x} \in [r]^n : w_1(\underline{x}) = \ldots = w_r(\underline{x}) = n/r\} .$$

It is clear that S does not contain a combinatorial line. At the same time, by Stirling's approximation, $\mu(S) \geq \Omega(1/n^{(r-1)/2})$ (where the constant in the $\Omega()$ notation depends on r) and therefore $\omega_r^{\mathrm{DHJ}}(n)$ cannot decrease exponentially.

By Theorem 1.8, $\omega_{\overline{Q}_r}(n)$ cannot decrease exponentially either. □

Better lower bounds for $\omega_r^{\mathrm{DHJ}}(n)$ are known, with the best ones established by the Polymath project [Pol10].

Theorem 4.1 ([Pol10], Theorem 1.3)**.** *Let $\ell \geq 1$ and $r := 2^{\ell-1} + 1$. There exists $C_\ell > 0$ such that for every $n \geq 2$ there exists a set $S \subseteq [r]^n$ with*

$$\mu(S) \geq \exp\left(-C_\ell \left(\log n\right)^{1/\ell}\right)$$

such that S does not contain a combinatorial line.

That is, for $r = 2^{\ell-1} + 1$, we have

$$\omega_{\overline{Q}_r}(n) \geq \exp\left(-C_r \left(\log n\right)^{1/\lceil \log r \rceil}\right) . \tag{4}$$

Inequality (4) is also interesting in the context of the two-prover parallel repetition lower bound by Feige and Verbitsky [FV02]. Recall that the upper bound of Raz (cf. (1)) exhibits a dependence on the answer set size. More specifically, it contains $1/\log\left|\overline{A}\right|$ term in the exponent. The example from [FV02] shows that if an exponential two-prover parallel repetition bound depends only on ϵ and $\left|\overline{A}\right|$, this term cannot be larger than $\log\log\left|\overline{A}\right| / \log\left|\overline{A}\right|$.

Our example implies that we can bring down this last term to $\left(\log\log\left|\overline{A}\right|\right)^{\epsilon} / \log\left|\overline{A}\right|$ for any $\epsilon > 0$, at the price of increasing the number of provers:

Theorem 4.2. *Let $\ell \geq 2$, $r := 2^{\ell-1} + 1$. There exists a constant $C_\ell > 0$ such that for each $n \geq 2$ there exists an r-prover game $\mathcal{G}$ with question set $\overline{Q}_r$, $\mathrm{val}(\mathcal{G}) \leq 1 - 1/r$ and an answer set $\overline{A}$ with size $\left|\overline{A}\right| \in [2^{rn}, 2^{2rn}]$ such that*

$$\mathrm{val}(\mathcal{G}^n) \geq \exp\left(-C_\ell n \cdot \frac{\left(\log\log\left|\overline{A}\right|\right)^{1/\ell}}{\log\left|\overline{A}\right|}\right) . \tag{5}$$

Proof. Fix ℓ and n and take the r-prover game $\mathcal{G}_S$ from Theorem 1.7 for the set $S \subseteq [r]^n$ from Theorem 4.1. One verifies that $\mathcal{G}_S$ has question set $\overline{Q}_r$ and that the answer alphabet size is $\left|\overline{A}\right| = (2^n \cdot n)^r \in [2^{rn}, 2^{2rn}]$.

Since S has no combinatorial line, we have $\mathrm{val}(\mathcal{G}_S) \leq 1 - 1/r$ and $\mathrm{val}(\mathcal{G}_S^n) \geq \mu(S) \geq \exp\left(-C_\ell \left(\log n\right)^{1/\ell}\right)$.

Noting that $n \geq \log\left|\overline{A}\right| / 2r$ and $\log n \leq \log\log\left|\overline{A}\right|$, we can establish (5):

$$\begin{aligned}\mathrm{val}(\mathcal{G}_S^n) &\geq \exp\left(-C_\ell \left(\log n\right)^{1/\ell}\right) = \exp\left(-C_\ell n \cdot \frac{\left(\log n\right)^{1/\ell}}{n}\right) \\ &\geq \exp\left(-C'_\ell n \cdot \frac{\left(\log\log\left|\overline{A}\right|\right)^{1/\ell}}{\log\left|\overline{A}\right|}\right) .\end{aligned}$$

□

As a final note, we reiterate that our lower bounds do *not* exclude the possibility of an "information theoretic" (see Section 1.6) parallel repetition bound. Furthermore, all results of this section concern games with at least three provers.

5 The Hales-Jewett Theorem and Coloring Games

As the name suggests, the density Hales-Jewett theorems is the density version of the earlier Hales-Jewett theorem [HJ63]. Inspired by Theorem 1.8, in this section we define *coloring games* and prove equivalence of their parallel repetition and the Hales-Jewett theorem. We are not aware of any previous works concerning coloring games.

5.1 The Hales-Jewett theorem

Definition 5.1. Let $r, c, n \in \mathbb{N}_{>0}$ and $C : [r]^n \to [c]$ be a coloring of $[r]^n$ with c colors. We say that there is a *monochromatic line* in C if there exists a combinatorial pattern $\underline{b}$ such that $C(\underline{b}(1)) = \ldots = C(\underline{b}(r))$.

For $r, n \in \mathbb{N}_{>0}$, let

$$\omega_r^{\mathrm{HJ}}(n) := \min \{c : \exists C : [r]^n \to [c] \text{ with no monochromatic lines}\} \ .$$

$\Diamond$

Theorem 5.2 (Hales-Jewett theorem)**.** *For every $r \geq 2$:*

$$\lim_{n\to\infty} \omega_r^{\mathrm{HJ}}(n) = +\infty \ .$$

Remark 5.3. Even though the Hales-Jewett theorem follows easily from the density version, better, primitive recursive bounds for $\omega_r^{\mathrm{HJ}}(n)$ are known [She88] compared to $\omega_r^{\mathrm{DHJ}}(n)$. $\Diamond$

5.2 Coloring games

Definition 5.4. An r-prover coloring game $\mathcal{G} = (\overline{Q}, \overline{A}, V)$ is given by a question set $\overline{Q} \subseteq Q^{(1)} \times \ldots \times Q^{(r)}$, an answer set $\overline{A} = A^{(1)} \times \ldots \times A^{(r)}$ and a function $V : \overline{Q} \times \overline{A} \to \mathcal{X}$ for some set $\mathcal{X}$.

The *color-value* of a game is

$$\mathrm{cVal}(\mathcal{G}) := \min_{\overline{\mathcal{S}}=(\mathcal{S}^{(1)},\ldots,\mathcal{S}^{(r)})} \left|V\left(\overline{Q}, \overline{\mathcal{S}}(\overline{Q})\right)\right| \ .$$

In other words, the provers are supposed to minimize the number of colors that the verifier can output instead of maximizing the probability of acceptance.

Given a coloring game $\mathcal{G}$, we define the parallel repetition similar to before. The only change is that the function V^n outputs a vector of n values which is obtained by applying V to every coordinate individually. $\Diamond$

Definition 5.5. For an r-prover question set $\overline{Q}$ let

$$\omega_{\overline{Q}}^{\mathrm{CPR}}(n) := \min_{\mathcal{G}} \mathrm{cVal}(\mathcal{G}^n) \ ,$$

where the minimum is over all coloring games $\mathcal{G}$ with question set $\overline{Q}$ and $\mathrm{cVal}(\mathcal{G}) \geq 2$. $\Diamond$

Theorem 5.6. *For every r-prover question set $\overline{Q}$:*

$$\lim_{n\to\infty} \omega_{\overline{Q}}^{\mathrm{CPR}}(n) = +\infty \ .$$

5.3 Equivalence of the theorems

Proof (Theorem 5.2 implies Theorem 5.6). In particular, we will show that

$$\omega_{\overline{Q}}^{\mathrm{CPR}}(n) \geq \omega_r^{\mathrm{HJ}}(n) \ , \tag{6}$$

where $r = \left|\overline{Q}\right|$.

Consider a coloring game $\mathcal{G}$ with the question set $\overline{Q}$ and suppose for a contradiction that there is a strategy for $\mathcal{G}^n$ where the verifier uses $c < \omega_r^{HJ}(n)$ colors. Fix such a strategy and identify the colors used in it with $[c]$. Consider now the map $C : \overline{Q}^n \to [c]$ which tells us what color the verifier will output in the repeated game for this strategy.

Theorem 5.2 implies that there is a pattern $\underline{b}$ such that $C(\underline{b}(1)) = \ldots = C(\underline{b}(r))$. This, however, implies that $\mathcal{G}$ has coloring value 1. To see this, consider the following prover strategy: on input $q^{(j)}$, prover j applies the repeated strategy with all $\star$ symbols in the pattern replaced with $q^{(j)}$, and for each other position she computes the question on the pattern input. Then, she responds with the response in the first star coordinate. □

Proof (Theorem 5.6 implies Theorem 5.2). In case of doubts the reader is advised to read the proof of Theorem 1.7 first. Recall the question set $\overline{Q}_r$ from Definition 1.6. In fact, we will prove that

$$\omega_r^{\mathrm{HJ}}(n) \geq \omega_{\overline{Q}_r}^{\mathrm{CPR}}(n) \ . \tag{7}$$

Let $c < \omega_{\overline{Q}_r}^{\mathrm{CPR}}(n)$ and fix an arbitrary function $C : [r]^n \to [c]$. We want to show that C has a monochromatic line. To this end, consider the following r-prover coloring game $\mathcal{G}_C$ over $\overline{Q}_r$:

The answer alphabet is $A^{(j)} := 2^{[n]} \times [n]$. After receiving answers $(T^{(j)}, z^{(j)})$, the verifier does checks similar as in the proof of Theorem 1.7: The sets $T^{(j)}$ should partition $[n]$ inducing a string $\underline{s} \in [r]^n$ and there should be a special coordinate $z = z^{(1)} = \ldots = z^{(r)}$ such that $z \in T^{(a)}$, where a is the special prover (i.e., the one that received 1).

If all the checks are passed, our verifier applies C on $\underline{s}$ and outputs the resulting element of $[c]$. Otherwise, the verifier outputs his question tuple $\overline{q}$ (which we assume to be not in $[c]$).

There is a strategy for $\mathcal{G}_C^n$ with coloring value c: letting $T^{(j)}$ be the set of coordinates in which prover j is special, prover j responds on coordinate i with $(T^{(j)}, i)$. Because $c < \omega_{\overline{Q}_r}^{\mathrm{CPR}}(n)$, we have $\mathrm{cVal}(\mathcal{G}_C) = 1$.

Consider a strategy for $\mathcal{G}_C$ that uses only one color. Since there are at least two questions, the verifier can never output $\overline{q}$. Consequently, $r \geq 3$ implies that there is a special coordinate z that the provers always output. Furthermore, the sets $T^{(j)}(0)$ are pairwise disjoint and the special prover a always outputs $T^{(a)}(0) \cup T(1)$, where $T(1) := [n] \setminus \left(T^{(1)}(0) \cup \ldots \cup T^{(r)}(0)\right)$ with $z \in T(1)$.

The coordinate sets $T^{(j)}(0)$ and $T(1)$ define a monochromatic combinatorial line. □

Figure 1: Doubling a bipartite graph. Fixed vertices in red, old vertices in green, new vertices in blue.

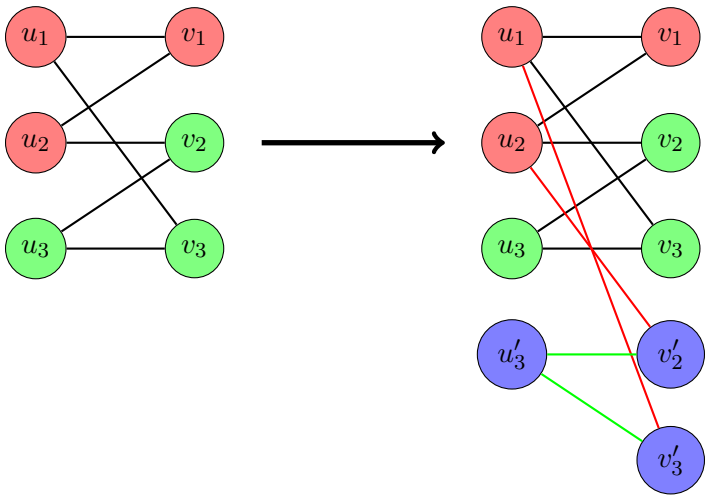

6 Constructability Implies Parallel Repetition

In this section we first define a class of constructible hypergraphs and then establish that all constructible question sets admit exponential parallel repetition. The main result of this section is Theorem 6.7.

6.1 Constructing hypergraphs by conditioning

We define constructability in the general case, but for intuition the reader is invited to think about bipartite graphs (i.e., $r = 2$). Recall our definitions of hypergraph homomorphisms (Definitions 2.4 and 2.5).

Definition 6.1. Given an r-regular, r-partite hypergraph $(Q^{(1)}, \ldots, Q^{(r)}, \overline{Q})$ and sets $P^{(j)} \subseteq Q^{(j)}$ we define its *section hypergraph* $(P^{(1)}, \ldots, P^{(r)}, \overline{P})$, where $\overline{P} \subseteq \overline{Q}$ consists of those hyperedges whose vertices are all in $P^{(1)} \cup \ldots \cup P^{(r)}$. $\Diamond$

In the graph case the section hypergraph corresponds to an induced subgraph.

Definition 6.2. Let $r \geq 2$. We recursively define the class of r-regular, r-partite hypergraphs that are *constructible by conditioning*:

1. A single hyperedge $\left(\{q^{(1)}\}, \ldots, \{q^{(r)}\}, \{(q^{(1)}, \ldots, q^{(r)})\}\right)$ is constructible.
2. If $(P^{(1)} \dot\cup Q^{(1)}, \ldots, P^{(r)} \dot\cup Q^{(r)}, \overline{P})$ is constructible, then $(P^{(1)} \dot\cup Q^{(1)} \dot\cup R^{(1)}, \ldots, P^{(r)} \dot\cup Q^{(r)} \dot\cup R^{(r)}, \overline{P} \cup \overline{Q})$ is constructible, where:

- $R^{(j)} := \{q' : q \in Q^{(j)}\}$ is a set of copies of vertices from $Q^{(j)}$. We say that the vertices in $P^{(j)}$ are *fixed*, vertices from $Q^{(j)}$ are *old* and vertices from $R^{(j)}$ are *new*.
- We say that a hyperedge is fixed if all of its vertices are fixed. For a hyperdge $\overline{q} \in \overline{P}$ that is not fixed we define $\overline{q}'$ as the hyperedge formed from $\overline{q}$ by replacing all of its old vertices by their respective copies.
  Then, $\overline{Q} := \big\{\overline{q}' : \overline{q} \in \overline{P}, \overline{q} \text{ is not fixed}\big\}$.

In this case we say that $\overline{P} \cup \overline{Q}$ was constructed from $\overline{P}$ by *doubling* $Q^{(1)} \cup \ldots \cup Q^{(r)}$.

Figure 1 can be consulted for an example in the graph case.

3. If $(P^{(1)} \dot\cup Q^{(1)}, \ldots, P^{(r)} \dot\cup Q^{(r)}, \overline{Q})$ is constructible and $(P^{(1)}, \ldots, P^{(r)}, \overline{P})$ is a section hypergraph of $\overline{Q}$ such that there exists a homomorphism from $\overline{Q}$ to $\overline{P}$ which is identity on $P^{(1)} \cup \ldots \cup P^{(r)}$, then $\overline{P}$ is constructible.

   In such case we say that $\overline{Q}$ *collapses onto* $\overline{P}$.

$\Diamond$

Observe that the doubling operation never produces hyperedges incident to both old and new vertices.

To give some intuition on the conditioning operations we state two simple properties.

Claim 6.3. *Let $r \geq 2$. Every r-partite hypergraph can be collapsed onto one of its hyperedges.*

Definition 6.4. The *complete* hypergraph on $Q^{(1)}, \ldots, Q^{(r)}$ is $(Q^{(1)}, \ldots, Q^{(r)}, Q^{(1)} \times \ldots \times Q^{(r)})$. $\Diamond$

Claim 6.5. *Let $r \geq 2$ and $Q^{(1)}, \ldots, Q^{(r)}$ be finite sets. The complete hypergraph on $Q^{(1)}, \ldots, Q^{(r)}$ is constructible.*

Proof. Let $k^{(j)} := |Q^{(j)}|$, $Q^{(j)} = \{q_1^{(j)}, \ldots, q_{k^{(j)}}^{(j)}\}$ and $P^{(j)} := \{q_1^{(j)}\}$. Start the construction with a single hyperedge $(P^{(1)}, \ldots, P^{(r)}, P^{(1)} \times \ldots \times P^{(r)})$.

The complete hypergraph is constructed in r stages. In the j-th stage vertex $q_1^{(j)}$ is doubled $k^{(j)} - 1$ times with all other vertices fixed. Observe that after the j-th stage the current hypergraph is the complete hypergraph on $Q^{(1)}, \ldots, Q^{(j)}, P^{(j+1)}, \ldots, P^{(r)}$. $\square$

Remark 6.6. As a matter of fact, if $|Q^{(1)}| = \ldots = |Q^{(r)}| = 2^k$, then it is not difficult to see that the complete hypergraph on $Q^{(1)}, \ldots, Q^{(r)}$ can be constructed with rk doublings. $\Diamond$

6.2 Constructability implies parallel repetition

Our goal in this section is the following quantitative version of Theorem 1.10:

Theorem 6.7. *Let $\overline{Q}$ be an r-prover question set that is constructible by conditioning using k doublings (and an arbitrary number of collapses). Let $M := \left|\overline{Q}\right|$. Then,*

$$\omega_{\overline{Q}}(n) \leq 3\exp\left(-n/M^{2^{k+1}}\right) .$$

In particular, $\overline{Q}$ admits exponential parallel repetition.

A very rough proof outline is as follows: first, exponential parallel repetition is equivalent to exponential decrease of the threshold for good homomorphism vectors (cf. Definition 3.4). Second, we show that the existence of good homomorphism vectors is implied by existence of a probability distribution over $\mathrm{Hom}\left(\overline{Q},\overline{Q}\right)$ with certain set hitting properties. Third, we prove that such distribution exists for every constructible $\overline{Q}$.

We elaborate those three steps in the next subsections.

6.2.1 Good question sets

For convenience we restate Definition 3.4:

Definition 3.4. Let $\overline{Q}$ be an r-prover question set and let $\underline{Q} := \overline{Q}^n$ be its n-fold parallel repetition. Let $S \subseteq \underline{Q}$ with $\mu(S) = |S|/\left|\overline{Q}\right|^n$ and let $\underline{f} = (f_1,\ldots,f_n)$ be a vector of n homomorphisms of $\overline{Q}$. We say that $\underline{f}$ is *good* for S if:

- For every $\overline{q} \in \overline{Q}$ we have that $\underline{f}(\overline{q}) := (f_1(\overline{q}),\ldots,f_n(\overline{q})) \in S$.
- There exists $i \in [n]$ such that f_i is identity.

We say that the question set $\overline{Q}$ is (n,ϵ)-good if for every $S \subseteq \underline{Q}$ with $\mu(S) \geq \epsilon$ there exists a vector of homomorphisms that is good for S. ◊

Observe that a vector of n identities $\underline{f} = (\mathrm{Id},\ldots,\mathrm{Id})$ is good for the whole space $\underline{Q}$ and hence every question set $\overline{Q}$ is $(n,1)$-good.

Remark 6.8. Note that if $\overline{Q}$ is (n,ϵ)-good, then it is also $(n+1,\epsilon)$-good. This is because given $S \subseteq \overline{Q}^{n+1}$ we can set f_{n+1} to be a constant homomorphism such that the (relative) measure $\mu(S)$ does not decrease conditioned on $\overline{q}_{n+1} = f_{n+1}(\overline{q})$. Then we can get $f_1,\ldots,f_n$ from the assumption that $\overline{Q}$ is (n,ϵ)-good. ◊

Definition 6.9. Let $\overline{Q}$ be a question set and $n \in \mathbb{N}_{>0}$. We define

$$\omega^{\mathrm{good}}_{\overline{Q}}(n) := \inf\left\{\epsilon : \overline{Q} \text{ is } (n,\epsilon)\text{-good}\right\} .$$

We say that $\overline{Q}$ is *good* if $\lim_{n\to\infty}\omega^{\mathrm{good}}_{\overline{Q}}(n) = 0$. ◊

The value of $\omega_{\overline{Q}}^{\text{good}}(n)$ is an upper bound on the parallel repetition rate $\omega_{\overline{Q}}(n)$:

Lemma 6.10. *Let $\overline{Q}$ be a question set. Then,*

$$\omega_{\overline{Q}}(n) \leq \omega_{\overline{Q}}^{\text{good}}(n) \ .$$

Proof. Assume otherwise, i.e., that there exists a game $\mathcal{G}$ with question set $\overline{Q}$ and $\text{val}(\mathcal{G}) < 1$ such that $\text{val}(\mathcal{G}^n) > \omega_{\overline{Q}}^{\text{good}}(n)$. We construct a perfect strategy for $\mathcal{G}$, which is a contradiction.

Fix an optimal strategy for $\mathcal{G}^n$ and let $S \subseteq \overline{Q}$ with $\mu(S) > \omega_{\overline{Q}}^{\text{good}}(n)$ be the set of question vectors in the repeated game for which the players win.

Let $\underline{f}$ be a vector of homomorphisms of $\overline{Q}$ that is good for S and let i be a coordinate where f_i is identity.

A strategy for the game $\mathcal{G}$ for the j-th prover is as follows: Given $q^{(j)} \in Q^{(j)}$, obtain $\underline{f}^{(j)}(q^{(j)}) = (f_1^{(j)}(q^{(j)}), \ldots, f_n^{(j)}(q^{(j)}))$. Then, consider the answer of the j-th prover on $\underline{f}^{(j)}(q^{(j)})$ in the strategy for $\mathcal{G}^n$. Finally, output the i-th coordinate of that answer.

Since for every $\overline{q} = (q^{(1)}, \ldots, q^{(r)}) \in \overline{Q}$ we have that $\underline{f}(\overline{q}) = (\underline{f}^{(1)}(q^{(1)}), \ldots, \underline{f}^{(r)}(q^{(r)})) \in S$, when applying the above strategy the provers are always winning on all coordinates of $\mathcal{G}^n$. Since $f_i(\overline{q}) = \overline{q}$, their answers on the i-th coordinate are winning for $\overline{q}$ in the game $\mathcal{G}$. Therefore, $\text{val}(\mathcal{G}) = 1$, a contradiction. □

Remark 6.11. Lemma 6.10 is essentially what is called in the literature the forbidden subgraph method. However, our formulation with homomorphisms is different than the one that uses forbidden subgraphs, e.g., in [FV02].

Verbitsky [Ver95, FV02] showed that the forbidden subgraph method is universal, i.e., for a connected question set $\overline{Q}$ there is $\omega_{\overline{Q}}(n) = \omega_{\overline{Q}}^{\text{good}}(n)$ (cf. Theorem 3.5). ◊

6.2.2 Proving that $\overline{Q}$ is good with probabilistic method

Lemma 6.12. *Let $\overline{Q}$ be a question set and let $\mathcal{H}$ be a distribution over* $\text{Hom}(\overline{Q}, \overline{Q})$ *such that:*

1. *If $\underline{f} = (f_1, \ldots, f_n)$ is sampled such that f_i is i.i.d. in $\mathcal{H}$, then:*

$$\forall S \subseteq \overline{Q} : \Pr\left[\forall \overline{q} \in \overline{Q} : \underline{f}(\overline{q}) \in S\right] \geq c(\mu(S)) \ ,$$

   where $c(\mu) > 0$ if $\mu > 0$.

2. *$\mathcal{H}(\text{Id}) > 0$.*

Then, $\overline{Q}$ is good. Furthermore, if $\mathcal{H}(\text{Id}) \geq \epsilon > 0$ and $c(\mu) \geq \mu^C / C$ for some $C \geq 1$, then $\omega_{\overline{Q}}^{\text{good}}(n) \leq 3\exp(-\epsilon n / C)$.

Proof. Let $\epsilon := \mathcal{H}(\mathrm{Id})$ and $\mu \in (0,1]$. For $S \subseteq \overline{Q}$ with $\mu(S) = \mu$, define the event

$$\mathcal{E} :\equiv \forall \overline{q} \in \overline{Q} : \underline{f}(\overline{q}) \in S \wedge \exists i \in [n] : f_i = \mathrm{Id} \ .$$

Since $\Pr[\forall \overline{q} \in \overline{Q} : \underline{f}(\overline{q}) \in S] \geq c(\mu)$ and $\Pr[\exists i : f_i = \mathrm{Id}] = 1 - (1-\epsilon)^n$, by union bound, if

$$\Pr[\forall i : f_i \neq \mathrm{Id}] = (1-\epsilon)^n \leq c(\mu)/2 \ , \tag{8}$$

then $\Pr[\mathcal{E}] > 0$. Therefore, if we choose n such that (8) holds, then $\overline{Q}$ is (n,μ)-good. Since for arbitrary $\mu \in (0,1]$ we found that $\overline{Q}$ is (n,μ)-good for n big enough, $\overline{Q}$ must be good.

Furthermore, if $c(\mu) \geq \mu^C/C$, setting

$$\mu := (2C(1-\epsilon)^n)^{1/C} \leq (2C)^{1/C} \cdot \exp(-\epsilon n/C) \leq 3\exp(-\epsilon n/C) \ ,$$

we see that:

$$(1-\epsilon)^n = \mu^C/2C \leq c(\mu)/2 \ ,$$

and therefore $\omega_{\overline{Q}}^{\mathrm{good}}(n) \leq 3\exp(-\epsilon n/C)$. □

6.2.3 Same-set hitting homomorphism spaces

Lemma 6.13. *Let $\overline{P}$ be an r-partite hypergraph constructible using k doublings (and an arbitrary number of collapses) and let $\overline{Q}$ be another r-partite hypergraph.*

Then, there exists a distribution $\mathcal{H}$ over $\mathrm{Hom}(\overline{P},\overline{Q})$ such that:

1. *If $\underline{f} = (f_1,\ldots,f_n)$ is sampled such that f_i is i.i.d. in $\mathcal{H}$, then:*

$$\forall S \subseteq \overline{Q} : \Pr\left[\forall \overline{p} \in \overline{P} : \underline{f}(\overline{p}) \in S\right] \geq \mu(S)^C \ ,$$

   where $C = 2^k$.

2. *$\min_{f \in \mathrm{Hom}(\overline{P},\overline{Q})} \mathcal{H}(f) \geq 1/M^C$, where $C = 2^k$ and $M = \left|\overline{Q}\right|$.*

This lemma is inspired by the paper [HHM15] in the following way: Let f be a random homomorphism sampled according to $\mathcal{H}$ and let $\overline{P} = \{\overline{p}^{(1)},\ldots,\overline{p}^{(k)}\}$. We can think of $\mathcal{H}$ as a k-step random process with the steps given by $f(\overline{p}^{(1)}),\ldots,f(\overline{p}^{(k)})$. Then, the first condition in Lemma 6.13 is equivalent to saying that $\mathcal{H}$ is polynomially same-set hitting as defined in [HHM15].

Later we will apply Lemma 6.13 with $\overline{P} = \overline{Q}$.

Proof. The proof proceeds by induction on the structure of $\overline{P}$. To achieve the constant C as claimed, we need to show the base case with $C = 1$ and then argue that a collapse preserves C and that a doubling increases C at most twice.

1. If $\overline{P}$ is a single hyperedge, then $\mathrm{Hom}(\overline{P}, \overline{Q})$ is isomorphic to $\overline{Q}$. Setting $\mathcal{H}(f_{\overline{q}}) := 1/M$ for $\overline{q} \in \overline{Q}$ one can easily see that both 1 and 2 are satisfied with $C = 1$.

2. Assume that $\overline{P}$ was constructed by doubling a hypergraph $\overline{P}_0$. Let A be the set of fixed vertices, B the old vertices and B' the new vertices (regardless of the player they belong to). Therefore the vertex set of $\overline{P}_0$ is $A \cup B$ and the vertex set of $\overline{P}$ is $A \cup B \cup B'$.

   We are going to write homomorphisms $f \in \mathrm{Hom}(\overline{P}_0, \overline{Q})$ as $f = (f_A, f_B)$ and $f \in \mathrm{Hom}(\overline{P}, \overline{Q})$ as $f = (f_A, f_B, f_{B'})$.

   Observe that

$$\begin{aligned}\mathrm{Hom}(\overline{P}, \overline{Q}) = \big\{ (f_A, f_B, f_{B'}) : (f_A, f_B) \in \mathrm{Hom}(\overline{P}_0, \overline{Q}) \\ \wedge\, (f_A, f_{B'}) \in \mathrm{Hom}(\overline{P}_0, \overline{Q}) \big\} \,, \end{aligned} \tag{9}$$

   where we abused the notation in the expression $(f_A, f_{B'})$: this is justified from the definition of the doubling operation.

   By induction, there exists a distribution $\mathcal{H}_0$ on $\mathrm{Hom}(\overline{P}_0, \overline{Q})$ satisfying 1 and 2 for some $C_0 > 0$. Let $H = (H_A, H_B)$ be a random variable distributed according to $\mathcal{H}_0$. Define:

$$\begin{aligned}\mathcal{H}(f_A, f_B, f_{B'}) := \Pr[H_A = f_A] \cdot \Pr[H_B = f_B \mid H_A = f_A] \\ \cdot \Pr[H_B = f_{B'} \mid H_A = f_A] \,. \end{aligned} \tag{10}$$

   By (9), (10) defines a probability distribution. Furthermore:

$$\mathcal{H}(f_A, f_B, f_{B'}) \geq \mathcal{H}_0(f_A, f_B) \cdot \mathcal{H}_0(f_A, f_{B'}) \geq \epsilon^{2C_0} \,.$$

   As for condition 1, let $\overline{E}_A$ be the fixed hyperedges of $\overline{P}_0$ (i.e., those that have all their vertices in A) and $\overline{E}_B$ and $\overline{E}_{B'}$ be the hyperedges of $\overline{P}$ that have vertices incident to B and B', respectively. Note that $\overline{E}_A$ and $\overline{E}_B$ form a partition of $\overline{P}_0$ and $\overline{E}_A$, $\overline{E}_B$ and $\overline{E}_{B'}$ form a partition of $\overline{P}$.

   Recall that $\underline{f} = (f_1, \ldots, f_n)$ is a random vector with coordinates sampled i.i.d. from $\mathcal{H}$. We are going to decompose $\underline{f} = (\underline{f}_A, \underline{f}_B, \underline{f}_{B'})$ in the natural way. Fix $S \subseteq \overline{Q}^n$ and define the event $\mathcal{E} :\equiv \forall \overline{p} \in \overline{E}_A : \underline{f}(\overline{p}) \in S$.

We estimate, using Jensen's inequality in (11):

$$\begin{aligned}
&\Pr\left[\forall \overline{p} \in \overline{P} : \underline{f}(\overline{p}) \in S\right] \\
&= \mathrm{E}\left[\mathrm{E}\left[\mathbb{1}_{\mathcal{E}} \cdot \Pr\left[\forall \overline{p} \in \overline{E}_B \cup \overline{E}_{B'} : \underline{f}(\overline{p}) \in S \mid \underline{f}_A\right] \mid \underline{f}_A\right]\right] \\
&= \mathrm{E}\left[\mathrm{E}\left[\mathbb{1}_{\mathcal{E}} \cdot \Pr\left[\forall \overline{p} \in \overline{E}_B : \underline{f}(\overline{p}) \in S \mid \underline{f}_A\right]^2 \mid \underline{f}_A\right]\right] \\
&= \mathrm{E}\left[\mathrm{E}\left[\left(\mathbb{1}_{\mathcal{E}} \cdot \Pr\left[\forall \overline{p} \in \overline{E}_B : \underline{f}(\overline{p}) \in S \mid \underline{f}_A\right]\right)^2 \mid \underline{f}_A\right]\right] \\
&= \mathrm{E}\left[\mathrm{E}\left[\mathbb{1}_{\mathcal{E}} \cdot \Pr\left[\forall \overline{p} \in \overline{E}_B : \underline{f}(\overline{p}) \in S \mid \underline{f}_A\right] \mid \underline{f}_A\right]^2\right] \\
&\geq \mathrm{E}\left[\mathrm{E}\left[\mathbb{1}_{\mathcal{E}} \cdot \Pr\left[\forall \overline{p} \in \overline{E}_B : \underline{f}(\overline{p}) \in S \mid \underline{f}_A\right] \mid \underline{f}_A\right]\right]^2 \qquad (11) \\
&= \Pr\left[\forall p \in P_0 : \underline{f}(p) \in S\right]^2 \geq \mu^{2C_0} \ .
\end{aligned}$$

3. The last case considers $\overline{P}$ constructed by collapsing some $\overline{P}_0$. Let A be the vertex set of $\overline{P}$ and $A \dot\cup B$ the vertex set of $\overline{P}_0$. Let $h \in \mathrm{Hom}(\overline{P}_0, \overline{P})$ be a homomorphism that defines this collapse.

   By induction, there exists a distribution $\mathcal{H}_0$ on $\mathrm{Hom}(\overline{P}_0, \overline{Q})$ satisfying properties 1 and 2 for some C_0. For $f \in \mathrm{Hom}(\overline{P}, \overline{Q})$, define

$$\mathcal{H}(f) := \sum_{\substack{g \in \mathrm{Hom}(\overline{P}_0, \overline{Q}) \\ g_A = f}} \mathcal{H}_0(g) \ .$$

   Since a restriction of a homomorphism is a homomorphism, $\mathcal{H}$ indeed is a probability distribution. Furthermore, since $\mathcal{H}(f) \geq \mathcal{H}(h \circ f) \geq \epsilon^{C_0}$, condition 2 is satisfied.

   Finally, let $\underline{f}_0$ be a vector of question homomorphisms sampled i.i.d. from $\mathcal{H}_0$ and recall that vector $\underline{f}$ is sampled i.i.d. from $\mathcal{H}$. To establish condition 1, we check that

$$\begin{aligned}
\Pr[\forall \overline{p} \in \overline{P} : \underline{f}(\overline{p}) \in S] &= \Pr[\forall \overline{p} \in \overline{P} : \underline{f}_0(\overline{p}) \in S] \\
&\geq \Pr[\forall \overline{p} \in \overline{P}_0 : \underline{f}_0(\overline{p}) \in S] \geq \mu^{C_0} \ .
\end{aligned}$$

□

6.2.4 Putting things together

Proof of Theorem 6.7. Let $\overline{Q}$ be an r-prover question set constructible by conditioning using k doublings with $\left|\overline{Q}\right| = M$.

By Lemma 6.13 applied for $\overline{P} = \overline{Q}$, Lemma 6.12 and Lemma 6.10,

$$\omega_{\overline{Q}}(n) \leq \omega_{\overline{Q}}^{\text{good}}(n) \leq 3\exp\left(-n/2^k M^{2^k}\right) \leq 3\exp\left(-n/M^{2^{k+1}}\right) .$$

Since $3\exp(-\alpha n) \leq \exp(-\alpha n/2)$ for n big enough, this implies that $\overline{Q}$ admits exponential parallel repetition. □

In particular, this recovers exponential parallel repetition for free games:

Corollary 6.14. *Let $\mathcal{G}$ be a free r-prover game with 2^k questions available to each prover. If $\mathcal{G}$ is non-trivial, then*

$$\text{val}(\mathcal{G}^n) \leq 3\exp(-n/M^{2M}) .$$

Proof. By Remark 6.6, the question set of game $\mathcal{G}$ can be constructed using rk doublings. The bound then follows from Theorem 6.7:

$$\text{val}(\mathcal{G}^n) \leq \omega_{\overline{Q}}(n) \leq 3\exp\left(-n/M^{2^{rk+1}}\right) = 3\exp\left(-n/M^{2M}\right) .$$

□

We note that quantitatively the bound for free games from Corollary 6.14 is much worse than the best known one by Feige [Fei91], which is $\exp\left(-\Omega\left(n/M\log M\right)\right)$.

7 Constructing Graphs with Treewidth Two

We turn to presenting the power of our system for proving parallel repetition. In particular, we show that all two-prover graphs with treewidth at most two are constructible.

Since in this section we deal only with two provers, we use more standard notation where a bipartite graph is denoted as $G = (X, Y, E)$. We will sometimes refer to vertices from X as "on the left" and from Y as "on the right".

Our main result here is Theorem 1.11.

7.1 Warm-up: forests are constructible

We start with showing that all forests are constructible, recovering the parallel repetition result by Verbitsky [Ver95]. We will later use Lemma 7.2 in the construction of series-parallel graphs.

Firstly, we note that it is only interesting to consider constructability of connected graphs (note that to create a new connected component one can double all vertices of an existing connected component):

Claim 7.1. *A bipartite graph G is constructible by conditioning if and only if all its connected components are constructible.*

We can always add a "fresh" leaf to a constructible graph G:

Lemma 7.2. *If $G = (X \dot\cup \{u\}, Y, E)$ is constructible, then $G' = (X \dot\cup \{u\}, Y \dot\cup \{v\}, E \cup \{(u, v)\})$ is also constructible.*

Proof. Pick an arbitrary edge (u, w) originating from u. Fix u and double all the other vertices. Then collapse all new vertices on the left onto u and all new vertices on the right onto w' (i.e., the copy of w). □

From Claim 7.1, iterated application of Lemma 7.2 and Theorem 6.7 we have:

Theorem 7.3. *Let G be a tree. Then, G is constructible by conditioning. In particular, if G is interpreted as a two-prover question set, then it admits exponential parallel repetition.*

7.2 Treewidth and series-parallel graphs

Definition 7.4 (Treewidth)**.** Let G be a simple graph. A *tree decomposition* of G is a tree T, where each node (also called a *bucket*) corresponds to a subset of the vertices of G, with the following properties:

- For each vertex v of G, the buckets in which v appears form a non-empty, connected subgraph of T.
- For each edge e of G, there exists a bucket that contains both endpoints of e.

The *width* of a tree decomposition of G is the size of the biggest bucket minus one. The *treewidth* of G denoted by $\mathrm{tw}(G)$ is the smallest possible width of a tree decomposition of G. ◊

We will not discuss treewidth here, referring the reader to any standard textbook on graph theory. We note that a connected graph has treewidth one if and only if it is a tree.

To characterise graphs with treewidth two, we need to introduce the notion of *generalized series-parallel graphs*.

Definition 7.5 (Series-parallel graphs)**.** Let $G = (X, Y, E)$ be a bipartite graph and $u, v \in X \cup Y$. We call a tuple (X, Y, E, u, v) an *oriented bipartite graph*. We call the vertex u the *top* and v the *bottom*.

We define the class of generalized bipartite series-parallel oriented (in short: series-parallel oriented) graphs recursively:

1. Let $G = (\{a\}, \{b\}, \{(a, b)\})$ be a single edge. Then, both (G, a, b) and (G, b, a) are series-parallel oriented graphs.

2. Let $G_1 = (X_1, Y_1, E_1, u, v)$ and $G_2 = (X_2, Y_2, E_2, v, w)$ be series-parallel oriented graphs such that $(X_1 \cup Y_1) \cap (X_2 \cup Y_2) = \{v\}$ and $v \in (X_1 \cap X_2) \cup (Y_1 \cap Y_2)$.

   Then, $G := (X_1 \cup X_2, Y_1 \cup Y_2, E_1 \cup E_2, u, w)$ is a series-parallel oriented graph.

   We say that G is a *series composition* of G_1 and G_2 with G_1 on top and G_2 at the bottom.

3. Let $G_1 = (X_1, Y_1, E_1, u, v)$ and $G_2 = (X_2, Y_2, E_2, v, w)$ be series-parallel oriented graphs satisfying the same preconditions as for the series composition.

   Then, both $G := (X_1 \cup X_2, Y_1 \cup Y_2, E_1 \cup E_2, u, v)$ and $G' := (X_1 \cup X_2, Y_1 \cup Y_2, E_1 \cup E_2, v, w)$ are series-parallel graphs.

   We say that G and G' are a *generalized series composition* of G_1 and G_2. We say that G_1 is the *primary graph* of G and that G_2 is the primary graph of G'.

4. Let $G_1 = (X_1, Y_1, E_1, u, v)$ and $G_2 = (X_2, Y_2, E_2, u, v)$ be series-parallel oriented graphs such that $(X_1 \cup Y_1) \cap (X_2 \cup Y_2) = \{u, v\}$ and $\{u, v\} \subseteq (X_1 \cap X_2) \cup (Y_1 \cap Y_2)$.

   Then, $G := (X_1 \cup X_2, Y_1 \cup Y_2, E_1 \cup E_2, u, v)$ is also a series-parallel oriented graph.

   We call G a *parallel composition* of G_1 and G_2.

We say that a bipartite graph G is series-parallel if there exist vertices u, v such that (G, u, v) is an oriented series-parallel graph. ♢

We refer the reader to Figure 2 for intuitive understanding of the composition operations.

The requirement that the vertices by which the bipartite graphs are joined belong to the set $(X_1 \cap X_2) \cup (Y_1 \cap Y_2)$ ensures that they belong to the same side of the graph and therefore the bipartedness is preserved. On the other hand, observe that the top and the bottom can lie either on the same or the opposite sides of the bipartite graph.

In the literature the (not necessarily bipartite) graphs constructed with series and parallel composition are usually called series-parallel, and graphs that are constructed also with generalized composition are called generalized series-parallel. Incidentally, a connected graph is generalized series-parallel if and only if all its biconnected components are series-parallel (see, e.g., [Bod07]).

From now on, by "series-parallel" we will always mean the generalized bipartite series-parallel graph from Definition 7.5.

We will use the following useful characterisation of graphs with treewidth at most two:

Theorem 7.6. *A connected bipartite graph G has treewidth at most two if and only if G is series-parallel.*

For a proof of Theorem 7.6 see [HHC99]. Their proof concerns the case of general (non-bipartite) graphs, but it is easy to see that a connected generalized series-parallel graph is bipartite if and only if it can be constructed with additional restrictions as in Definition 7.5.

Figure 2: Illustration of the composition operations. The spines are drawn in continuous red. Note that $\mathcal{S}(G_2)$ is not part of $\mathcal{S}(G)$ in cases of generalized and parallel composition, hence the red dashed line.

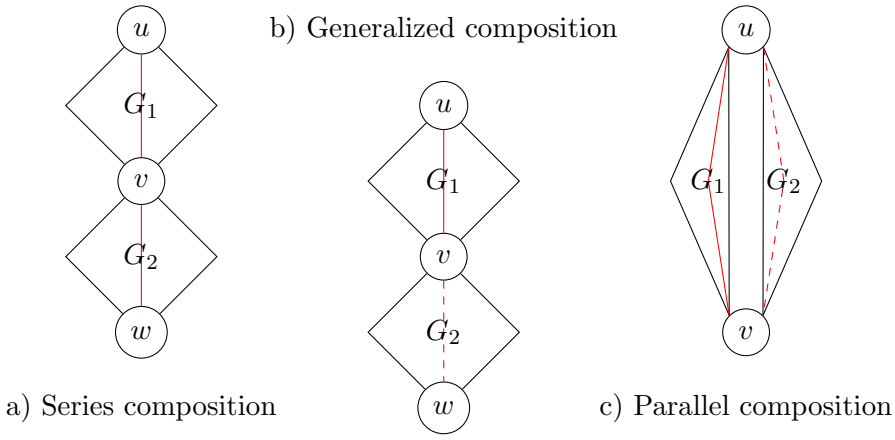

7.3 Generalized series-parallel construction

Recall that the main theorem of this section is:

Theorem 1.11. *Every bipartite graph G with treewidth at most two is constructible by conditioning. In particular, if G is interpreted as a two-prover question set, then it admits exponential parallel repetition.*

Due to Theorem 6.7, Claim 7.1 and Theorem 7.6, to establish Theorem 1.11 it is enough to show that series-parallel graphs are constructible. We spend the rest of this section to achieve that goal.

Definition 7.7. Let G be an oriented series-parallel graph. We define its (not oriented) subgraph $\mathcal{S}(G)$ and call it its *spine*. The definition follows the recursive pattern of Definition 7.5:

1. If G is a single edge, its spine is the whole of G.

2. If G is a series composition of G_1 and G_2, then $\mathcal{S}(G)$ consists of $\mathcal{S}(G_1)$ and $\mathcal{S}(G_2)$ taken together.

3. If G is a generalized composition of G_1 and G_2 with G_1 as the primary graph, then $\mathcal{S}(G)$ is equal to $\mathcal{S}(G_1)$.

4. If G is a parallel composition of G_1 and G_2 and $\mathcal{S}(G_1)$ has no more edges than $\mathcal{S}(G_2)$, then $\mathcal{S}(G)$ is equal to $\mathcal{S}(G_1)$. Otherwise, it is equal to $\mathcal{S}(G_2)$.

$\Diamond$

Observe that the spine is always an induced path between the top and the bottom of G. As a matter of fact, it is a shortest path from top to bottom in G. Furthermore, the length of the spine $L(G)$ is given as:

1. One, if G is a single edge.

2. $L(G_1) + L(G_2)$, if G is a series composition of G_1 and G_2.

3. $L(G_1)$, if G is a generalized composition of G_1 and G_2 with G_1 as the primary graph.

4. $\min(L(G_1), L(G_2))$, if G is a parallel composition of G_1 and G_2.

Finally, note that if G is a parallel composition of G_1 and G_2, then due to the bipartedness $L(G_1)$ and $L(G_2)$ must have the same parity.

Recall the graph construction operations from Definition 6.2. A series-parallel graph can always be collapsed onto its spine:

Lemma 7.8. *Let G be an oriented series-parallel graph. Then, G (treated as an unoriented graph) can be collapsed onto its spine.*

Proof. By induction on the series-parallel structure of G. If G is a single edge, it is clear. If G is a series composition of G_1 and G_2, then by induction G_1 and G_2 can be collapsed onto their respective spines.

If G is a generalized composition of G_1 and G_2, assume w.l.o.g. that G_1 is the primary graph and let v be the bottom vertex of G_1. Then, by induction, G_1 can be collapsed onto its spine $\mathcal{S}(G_1) = \mathcal{S}(G)$. On the other hand, all of G_2 can be collapsed onto the edge (v, w), where w is the neighbor of v in $\mathcal{S}(G_1)$.

In case (G, u, v) is a parallel composition of G_1 and G_2, assume w.l.o.g. that $\mathcal{S}(G_1)$ is not longer than $\mathcal{S}(G_2)$. Firstly, observe that the spine $\mathcal{S}(G_2)$ can be collapsed onto $\mathcal{S}(G_1) = \mathcal{S}(G)$: indeed, if we write the vertices of $\mathcal{S}(G_1)$ top-bottom as $(u_0 = u, u_1, \ldots, u_k = v)$ and analogously $\mathcal{S}(G_2)$ as $(v_0 = u, v_1, \ldots, v_{k+2\ell} = v)$, then the mapping:

$$f(u_i) := u_i$$
$$f(v_i) := \begin{cases} u_i & \text{if } i \leq k, \\ u_{k-(j \bmod 2)} & \text{if } i = k + j, \end{cases}$$

is a required homomorphism.

Finally, since by induction G_1 and G_2 can be collapsed onto their spines, and since the composition of homomorphisms is a homomorphism, G can be collapsed onto its spine. □

Recall that our objective is showing that every series-parallel graph is constructible.

Lemma 7.9. *Let (G, u, v) be an oriented series-parallel graph. Then, the spine $\mathcal{S}(G)$ can be extended to G using the doubling and collapsing operations. Furthermore, the construction preserves the following invariant:*

- *In every doubling step, the doubled vertices on the spine form its contiguous (possibly empty) subsegment.*

Since the spine of G can be constructed by repeated application of Lemma 7.2, Lemma 7.9 implies what we want. In the remainder we prove Lemma 7.9 after establishing a couple of technical preliminaries.

Remark 7.10. In the proof of Lemma 7.9 we will use the fact that whenever G is a composition of G_1 and G_2, the edges of G_1 and G_2 are disjoint.

This is not true if G is a parallel composition and there exists a direct edge from top to bottom in both G_1 and G_2, but any series-parallel G can be constructed without using this special case. ◊

Claim 7.11. *Let (G, u, v) be an oriented series-parallel graph which is a parallel composition. Then, there exists a series-parallel construction of (G, u, v) such that its final step is a parallel composition of G_1 and G_2 with the following properties:*

- *$L(G_1) \leq L(G_2)$.*
- *G_2 is a series composition.*

Proof. Firstly, note that whenever (G, u, v) is a generalized composition where the primary graph is a series or parallel composition, the order of those two compositions can be reversed without changing the final graph. Therefore, we can assume w.l.o.g. that whenever a graph is a generalized composition, its primary graph has spine of length one.

Let G be a parallel composition of H_1' and H_2'. If any of H_1' or H_2' is a parallel composition, recursively decompose them further until we are left with a collection of graphs $H_1, \ldots, H_k$ which are all series or generalized compositions or single edges.

Note that if we compose in parallel $H_1, \ldots, H_k$ in an arbitrary order, the end result will always be G.

Therefore, we can set G_2 as H_i with the longest spine and the parallel composition of the remaining H_i graphs as G_1. Due to Remark 7.10, the spine of G_2 must be longer than one, and therefore G_2 must be a series composition. □

Figure 3 illustrates the content of Claim 7.11.

Figure 3: The continuous red line is the spine of G. The dashed red line is the spine of G_2.

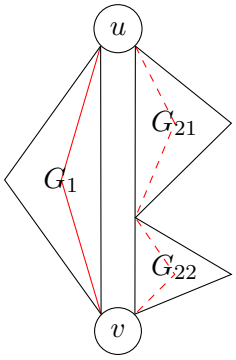

Proof of Lemma 7.9. Let (G, u, v) be an oriented series-parallel graph. We apply induction on the number of vertices of G and, secondarily, (in reverse) on the length of its spine.

1. If G is a single edge, there is nothing to prove (since $\mathcal{S}(G) = G$).

2. Assume that (G, u, v) is a series composition of (G_1, u, w) and (G_2, w, v). Recall that we need to extend $\mathcal{S}(G)$ to G. We do it in two stages, first extending $\mathcal{S}(G_1)$ to G_1 and then extending $\mathcal{S}(G_2)$ to G_2.

   By induction, we know how to extend $\mathcal{S}(G_1)$ to G_1. Now we will adapt this sequence of operations to the fact that also $\mathcal{S}(G_2)$ is present in the graph. We do it as follows:

   - Leave all the collapsing operations as they are (it is always possible to collapse onto a bigger graph).
   - For doubling operations that keep the vertex w fixed, keep all of $\mathcal{S}(G_2)$ fixed.
   - Finally, let us handle the doubling operations that double the vertex w. Let x be the neighbour of w on the spine $\mathcal{S}(G_1)$ and let y be x in case x is fixed and x' in case x is doubled. Note that the edge (y, w') is present in G_1 after doubling. To emulate this operation in G we double all of $\mathcal{S}(G_2)$ together with w and then collapse the new copy of $\mathcal{S}(G_2)$ onto the edge (y, w').

   Consult Figure 4 for the illustration of one of the cases.

   It is easy to see that as a result of this emulation we extend $\mathcal{S}(G)$ to a series composition of G_1 and $\mathcal{S}(G_2)$.

   Now we proceed in the same way to extend $\mathcal{S}(G_2)$ to G_2. The only difference is that in case w is doubled we need to double and collapse all of G_1 instead of just $\mathcal{S}(G_1)$. This does not pose a problem though, since G_1 can be collapsed onto $\mathcal{S}(G_1)$ which then can be collapsed as previously.

Figure 4: Handling series decomposition in case w and x are doubled.

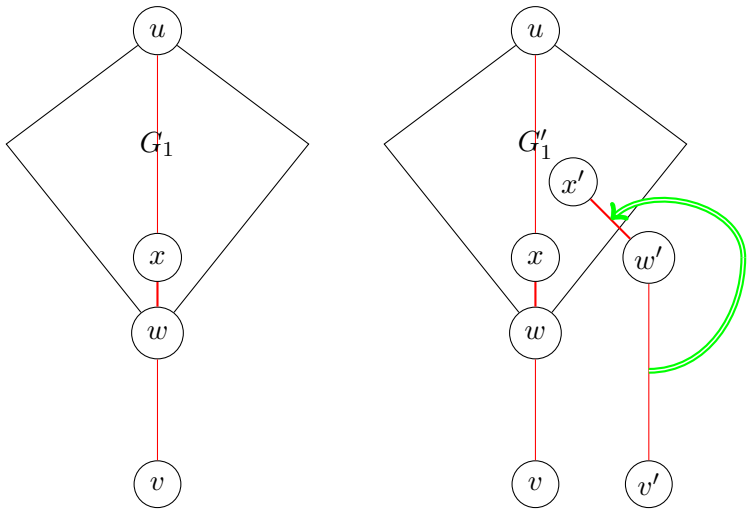

Finally, one easily checks that the "contiguous subsegment" invariant of Lemma 7.9 is preserved in this construction.

3. If (G, u, v) is a generalized composition, assume w.l.o.g. that it is a composition of the primary graph (G_1, u, v) and (G_2, v, w). Using Lemma 7.2 we can extend $\mathcal{S}(G_1)$ to $\mathcal{S}(G_1) \cup \mathcal{S}(G_2)$ and then proceed as in the series composition case.

4. Assume that (G, u, v) is a parallel composition of (G_1, u, v) and (G_2, u, v). By Claim 7.11, we can also assume that G_2 is a series composition of (G_3, u, w) and (G_4, w, v) and that $L(G_1) \leq L(G_2)$. In this point we address a subcase where additionally:

$$L(G_1) + L(G_3) < L(G_4) \ . \tag{12}$$

(G, u, v) is the parallel composition of (G_1, u, v) and the series composition of (G_3, u, w) and (G_4, w, v). Observe that we can also obtain (G, w, v) as the parallel composition of (G_4, w, v) and the series composition of (G_3, w, u) and (G_1, u, w). This is illustrated in Figure 5. Furthermore, due to (12) we have that $L(G, w, v) = L(G_3) + L(G_1) > L(G_1) = L(G, u, v)$.

To extend $\mathcal{S}(G, u, v) = \mathcal{S}(G_1)$ we proceed as follows: first, add $\mathcal{S}(G_3)$ on top of $\mathcal{S}(G_1)$ using Lemma 7.2. Then, extend $\mathcal{S}(G_3) \cup \mathcal{S}(G_1) = \mathcal{S}(G, w, v)$ to G using induction (which is applicable since the length of the spine increased).

Again, one easily checks that the contiguous subsegment invariant is preserved in this construction.

Figure 5: Rotating G which is a parallel composition.

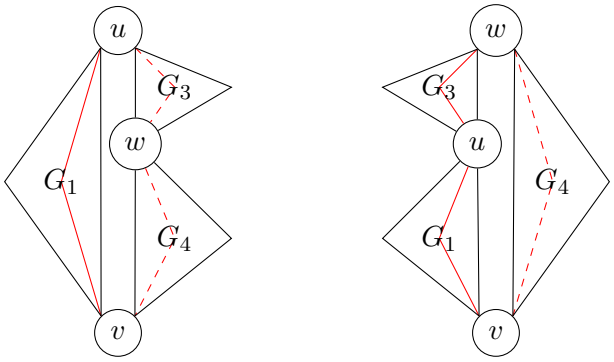

5. If G is a parallel composition and $L(G_1)+L(G_4) < L(G_3)$, we proceed symmetrically as in case 4.

6. Finally, let (G, u, v) be a parallel composition and:

$$L(G_1) + L(G_3) \geq L(G_4)\ , \tag{13}$$
$$L(G_1) + L(G_4) \geq L(G_3)\ . \tag{14}$$

   Again, we proceed in stages successively building (G_1, u, v), (G_3, u, w) and (G_4, w, v) using induction.

   We start with $\mathcal{S}(G) = \mathcal{S}(G_1)$, which we need to extend to G. First, by induction we extend $\mathcal{S}(G_1)$ to G_1. Next, we add $\mathcal{S}(G_2) = \mathcal{S}(G_3) \cup \mathcal{S}(G_4)$ as follows: let $a := L(G_1)$ and $b := L(G_2)$. Recall that $b \geq a$ and that a and b have the same parity.

   Using Lemma 7.2, add a path of length $(b-a)/2$ starting from u and let x be the endpoint of this path. Fix v and x and double all the other vertices. Finally, collapse the resulting copy of G_1 onto the path from v to u'.

   In the next stage, we work with the sequence that extends $\mathcal{S}(G_3)$ to G_3. We need to adapt it to additional edges we have in the graph. This is done as follows:

   - All collapsing operations stay the same.
   - Doubling operations that keep both u and w fixed fix all the vertices of G_1 and $\mathcal{S}(G_4)$.
   - In case at least one of u and w is doubled the arguments are very similar to each other. Therefore we present only the one where u is doubled and w is fixed. See Figure 6 for a graphical illustration.

     By inductive assumption, we know that a contiguous subpath of the spine $\mathcal{S}(G_3)$ is doubled. Assume that its doubled vertices go from u to y and the fixed ones from w to x (i.e., x and y are neighbours on the spine).

Figure 6: Handling parallel decomposition when segment from u to y is doubled. For clarity, G_1 and G_4 are drawn as spines only. The blue path is collapsed onto the green path.

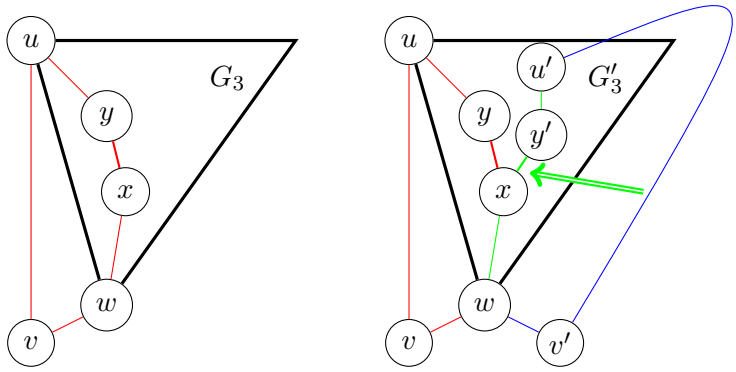

To emulate this case in G, double all vertices of G_1 and $\mathcal{S}(G_4)$ except of w. Next, collapse the new copy of G_1 onto its spine. Finally, collapse the resulting path $\mathcal{P}_1 := u' - v' - w$ onto the copy of $\mathcal{S}(G_3)$, i.e., $\mathcal{P}_2 := u' - y' - x - w$. This is possible due to (14): since the path $\mathcal{P}_1$ is at least as long as $\mathcal{P}_2$, $\mathcal{P}_1$ can be collapsed onto $\mathcal{P}_2$ as in the proof of Lemma 7.8.

Finally, we construct G_4 from $\mathcal{S}(G_4)$ in a very similar way. The only differences are that when emulating doubling we need to perform an additional collapse of G_3 onto $\mathcal{S}(G_3)$ and that we rely on the inequality (13) for the final collapse.

Again, one checks that the contiguous subsegment invariant is preserved throughout the whole process.

□

8 Some Graphs Are Not Constructible

It is an open question if all two-prover question sets admit exponential parallel repetition. One way to prove that they do would be to show that all graphs are constructible by conditioning. However, in this section we show that that is not the case, hence another way must be found to resolve this open question:

Definition 8.1. Let $n \in \mathbb{N}$ be even and greater or equal to 8. We define the *cycle with shortcuts* $\mathfrak{C}_n$ as the following simple graph: $V(\mathfrak{C}_n) := \{0, \dots, n-1\}$ and $\{u, v\} \in E(\mathfrak{C}_n)$ if and only if $|u - v| \in \{1, 3, n-3, n-1\}$. ◊

See Figure 7 for a drawing of $\mathfrak{C}_{12}$. Observe that $\mathfrak{C}_n$ is bipartite. We show:

Figure 7: A drawing of $\mathfrak{C}_{12}$.

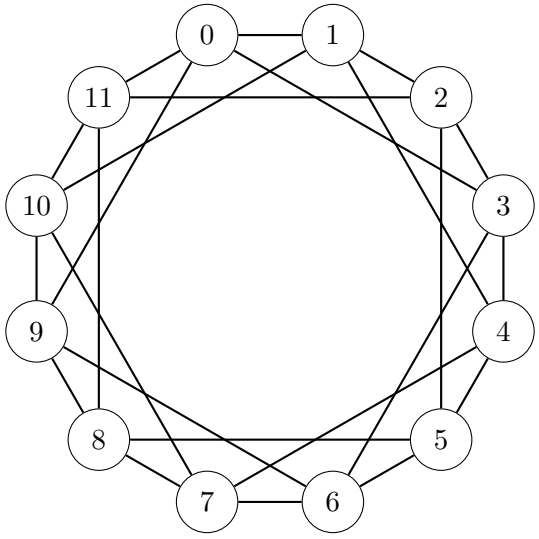

Theorem 8.2 (cf. Theorem 1.12)**.** *The cycle with shortcuts $\mathfrak{C}_{12}$ is not constructible by conditioning.*

Since any bipartite graph G joined with $\mathfrak{C}_{12}$ by a single vertex can be collapsed onto $\mathfrak{C}_{12}$, Theorem 8.2 implies the existence of an infinite family of graphs that are not constructible.

Our proof of Theorem 8.2 turns out to be somewhat involved and computer-assisted. Before we proceed with it, we explain why another natural proof idea fails.

8.1 "Warm-up": constructing all induced subgraphs

A natural idea to prove Theorem 8.2 would be to show for a certain graph G that if it is not already present as an induced subgraph in another graph H, then no doubling of H can produce an induced instance of G. It turns out that this approach must fail, since for every bipartite graph G we can construct a graph H such that G is an induced subgraph of H.

Definition 8.3. Let $k \geq 1$. We define the *set graph* $\mathfrak{S}_k := (X, Y, E)$ as follows:

- $X := [k]$.
- $Y := \{S \subset [k] : S \neq \emptyset\}$.
- $E := \{(x, S) : x \in S\}$.

$\Diamond$

Theorem 8.4. *The set graph $\mathfrak{S}_k$ is constructible by conditioning with $2(k-1)$ doublings.*

Proof. The proof is by induction on k. The graph $\mathfrak{S}_1$ is just a single edge. To construct $\mathfrak{S}_{k+1}$, start with constructing $\mathfrak{S}_k$ with $2(k-1)$ doublings.

We make a preliminary point to avoid confusion. Note that the right hand-side vertices of $\mathfrak{S}_k$ are labeled with subsets of $[k]$ such that for a vertex labeled with S we have that its neighborhood is equal to its label: $N(S) = S$. We will now perform some doublings and label the new vertices with subsets that contain $k+1$. However, for a new vertex with a label S it is not evident that $N(S) = S$: this is what we have to prove.

After constructing $\mathfrak{S}_k$, perform a doubling as follows: double all vertices labeled with S such that $k \in S$ and label each new vertex as $S \cup \{k+1\}$.

Then, perform a second doubling: double k and, again, all vertices labeled with S such that $k \in S$ and $k+1 \notin S$. This time label the copy of k as $k+1$ and a copy of S as $S \setminus \{k\} \cup \{k+1\}$.

Note that after the doublings $Y = \{S \subseteq [k+1] : S \neq \emptyset\}$. For $S \in Y$ let $N(S) := \{x \in X : (x, S) \in E\}$ be the neighborhood of S. We need to check that $N(S) = S$ for every label S. This holds by the following case analysis:

- Each vertex labeled with S such that $k+1 \notin S$ existed before the first doubling and its neighborhood did not change (since it was doubled in the second doubling in case $k \in S$).
- Each vertex labeled with S such that $\{k, k+1\} \subseteq S$ was created in the first doubling, at which point we had $N(S) = S \setminus \{k+1\}$. Then, it was fixed in the second doubling and $k+1$ was added to its neighborhood.
- Each vertex labeled with S such that $k \notin S$ and $k+1 \in S$ was created in the second doubling with $N(S) = S$.

Therefore, we can construct $\mathfrak{S}_{k+1}$ from $\mathfrak{S}_k$ in 2 doublings and $\mathfrak{S}_{k+1}$ from $\mathfrak{S}_1$ in $2k$ doublings. □

Remark 8.5. A modification of this construction can be used to construct $\mathfrak{S}_{k,r}$ with $X := [k]$, $Y := \{S \subseteq [k] : |S| = r\}$ and $E := \{(x, S) : x \in S\}$. ◊

Now we turn to the proof of Theorem 8.2.

8.2 Decomposing last two steps

Definition 8.6. Let u, v be two vertices arising during a construction of a bipartite graph G. We write $u \sim v$ if u and v are adjacent. For two sets of vertices A, B, we write $E(A, B)$ for the set of edges between A and B. We also write $G(A)$ for the graph induced by vertices in A. ◊

Note that the operators $\sim$, $E(\cdot,\cdot)$ and $G(\cdot)$ do not depend on the stage of the construction: doubling and collapsing only add and remove vertices, without changing existing adjacencies.

Lemma 8.7. *Let G be bipartite graph. If G is constructible, then it is constructible such that all the operations except for the last one are doublings.*

Proof. First, assume that in a construction of G there is a collapse operation immediately followed by a doubling operation. Assume that A is the set of the vertices collapsed in the first operation, B is the set of vertices that are fixed in the first operation and doubled onto B' in the second operation and C the set of vertices that are fixed throughout both operations (see Figure 8).

Figure 8: Transposing a collapse and a doubling.

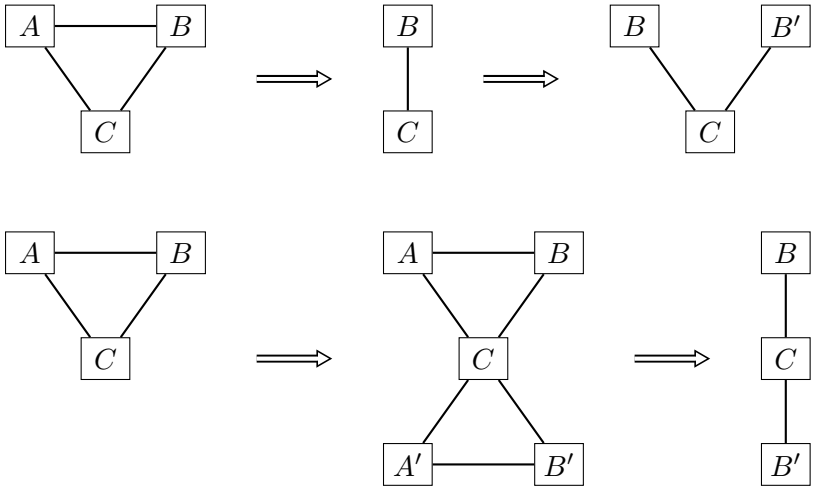

Then, those two operations can be exchanged as follows. First, double A onto A' and B onto B'. Then, collapse A onto $B \cup C$ and A' onto $B' \cup C$ (again see Figure 8). In both cases we end up with the same graph on vertices $B \cup B' \cup C$.

Finally, note that once all collapses are at the end of the sequence of the operations, they can be merged into a single collapse. □

Definition 8.8. We say that a graph G is *collapsible* onto a graph H, if H can be constructed from G by a single collapse operation. ◊

Lemma 8.9. *Let H be a constructible graph with at least two edges. There exists a construction of H such that:*

1. *The last operation is a collapse.*

2. All other operations are doublings.

3. Leting H_0 be the graph before the last doubling, H_0 is not collapsible onto H.

Proof. By Lemma 8.7, there exists a construction of H satisfying the first two conditions. Let us take such a construction with the smallest possible number of doublings. Since H is not a single edge, the number of doublings must be at least one.

If in this construction H_0 is collapsible onto H, the last doubling and the collapse can be replaced with a single collapse, which is a contradiction. □

Due to Lemma 8.7, we can assume that if the graph $\mathfrak{C}_{12}$ is constructible, the last two steps of its construction are, respectively, doubling and collapsing. Let us now divide the vertices of the construction depending on what happens to them in those last two steps (see Figure 9).

Figure 9: The last two steps in a construction of $\mathfrak{C}_{12}$.

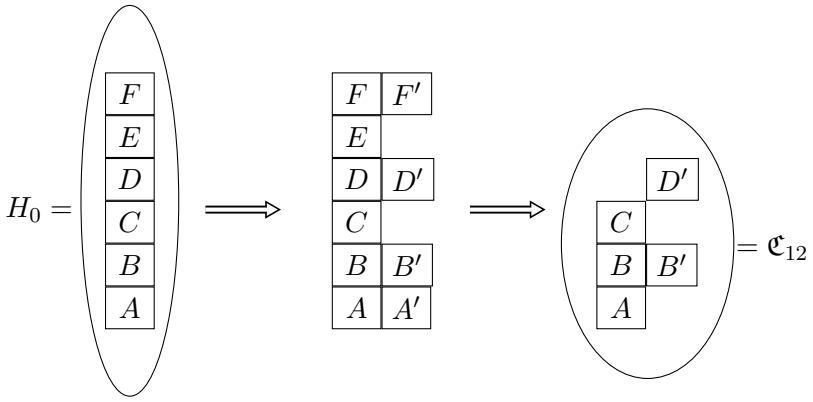

The division is as follows: A are vertices that are doubled onto A' in the first step, with A fixed and A' collapsed in the second step. B are vertices doubled onto B' in the first step with both B and B' fixed in the second step. C are vertices that are fixed throughout both steps. D are vertices doubled onto D' in the first step with D collapsed and D' fixed in the second step. E are vertices fixed in the first step and collapsed in the second step. Finally, F are vertices that are doubled onto F' in the first step with both F and F' collapsed in the second step.

One checks that this division covers all possible events in the last two steps. The final graph $\mathfrak{C}_{12}$ consists of vertices $A \cup B \cup B' \cup C \cup D'$.

Our proof of Theorem 8.2 goes as follows: First, we show that if the last two steps of a construction of $\mathfrak{C}_{12}$ are as above, it must be $B = \emptyset$ and $E(A, D) = \emptyset$. Then, we prove that if $B = \emptyset$ and $E(A, D) = \emptyset$, then the initial graph H_0 must have been collapsible onto $\mathfrak{C}_{12}$ in the first place. Parts of the proof are computer-assisted, with the codes of C++ programs provided in Appendix A.

8.3 Non-collapsible graphs never produce $\mathfrak{C}_{12}$

Lemma 8.10. *Let $\mathfrak{C}_{12}$ be constructed in two steps from some bipartite H_0, as above. It cannot be that $E = F = \emptyset$, $E(A, D) = \emptyset$ and $B \neq \emptyset$.*

Proof. Computer-assisted (enumerate all partitions of $\mathfrak{C}_{12}$ into $A \cup B \cup B' \cup C \cup D'$ together with a bijection between B and B', since $E(A, D) = \emptyset$ such a partition implies a unique $H_0 = G(A \cup B \cup C \cup D)$), see the program `non_empty_b.cpp` in Listing 2. □

Lemma 8.11. *Let $\mathfrak{C}_{12}$ be constructed in two steps from some bipartite H_0, as above. It cannot be that $B = E = F = \emptyset$ and $|E(A, D)| = 1$.*

Proof. Computer-assisted (enumerate all partitions of $\mathfrak{C}_{12}$ into $A \cup C \cup D'$ and all edges between A and D, again this implies a unique $H_0 = G(A \cup C \cup D)$), see the program `non_empty_ad.cpp` in Listing 3. □

Lemma 8.12. *Let $\mathfrak{C}_{12}$ be constructed in two steps from some bipartite H_0, as above. Then, it must be that $B = \emptyset$ and $E(A, D) = \emptyset$.*

Proof. Assume by contradiction that there exists a construction of $\mathfrak{C}_{12}$ with $B \neq \emptyset$ or $E(A, D) \neq \emptyset$.

Firstly, note that the same construction but with the vertices from $E \cup F$ deleted from the initial graph H_0 is valid and also results in $\mathfrak{C}_{12}$. Therefore, we can assume w.l.o.g. that $E = F = \emptyset$.

We now proceed in two cases. If $B \neq \emptyset$, we can additionally assume that $E(A, D) = \emptyset$. This is again due to the fact that if we deleted $E(A, D)$ edges from H_0, we would still obtain a valid construction that results in $\mathfrak{C}_{12}$ (cf. Figure 10). But $B \neq \emptyset$ and $E(A, D) = \emptyset$ is impossible due to Lemma 8.10.

On the other hand, assume that $B = \emptyset$ and $E(A, D) \neq \emptyset$. Then, by the same argument as before, we can also assume that the size of $E(A, D)$ is as small as possible, namely $|E(A, D)| = 1$ (cf. Figure 11). But this also yields a contradiction by Lemma 8.11. □

Figure 10: An illustration of Lemma 8.12, case $B \neq \emptyset$.

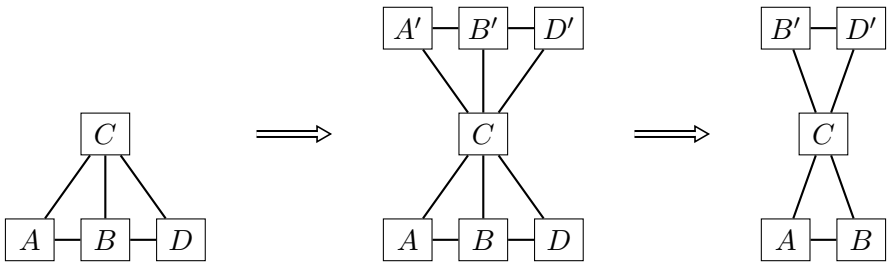

We need some additional concepts to deal with the remaining case $B = \emptyset$, $E(A, D) = \emptyset$.

Figure 11: An illustration of Lemma 8.12, case $B = \emptyset$, $|E(A, D)| = 1$. The edge between A and D is marked red.

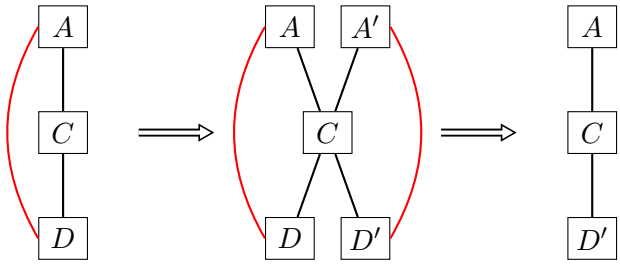

Definition 8.13. Let A and B be disjoint sets of vertices that exist at some point during a construction of a bipartite graph H. Assume that a doubling operation is performed and that all vertices from B (possibly together with some vertices from A and outside of $A \cup B$) are doubled.

Let B' be the set of copies of vertices from B. There is an obvious bijection between B and B' which we call the *natural bijection*. Similarly, we say that there is natural bijection between $A \cup B$ and $A \cup B'$. If this bijection is also an isomorphism between $G(A \cup B)$ and $G(A \cup B')$, we say that $G(A \cup B)$ and $G(A \cup B')$ are *naturally isomorphic*. ◊

Definition 8.14. Let $\mathfrak{C}_{12}$ be constructed from some H_0 in two steps, as above. We say that A' was *naturally collapsed* onto A, if in the collapse step each vertex of A' was collapsed onto A via the natural bijection. Analogously, we say that D was naturally collapsed onto D'. ◊

Lemma 8.15. *Let $\mathfrak{C}_{12}$ be constructed from some H_0 with one doubling and one collapse, as above. If $B = E = F = \emptyset$ and $E(A, D) = \emptyset$, then in the subsequent collapse either A' is naturally collapsed onto A or D is naturally collapsed onto D'.*

Proof. Computer-assisted (enumerate all partitions of $\mathfrak{C}_{12}$ into $A \cup C \cup D'$ and all possible collapses), see the program `natural_collapse.cpp` in Listing 4. □

Lemma 8.16. *Let $\mathfrak{C}_{12}$ be constructed from some H_0 with one doubling and one collapse, as above. If $B = \emptyset$ and $E(A, D) = \emptyset$, then in the subsequent collapse either A' is naturally collapsed onto A or D is naturally collapsed onto D'.*

Proof. Assume there exists a construction of $\mathfrak{C}_{12}$ from some $H = G(A \cup C \cup D \cup E \cup F)$ such that:

1. $B = \emptyset$ and $E(A, D) = \emptyset$.
2. A' does not naturally collapse onto A and D does not naturally collapse onto D'.

Then, the same construction with vertices $E \cup F$ omitted from H_0 is also valid and satisfies both conditions. Therefore, we can assume w.l.o.g. that $E = F = \emptyset$. Then, the result follows from Lemma 8.15. □

Figure 12: An illustration of Lemma 8.16, case when $E = F = \emptyset$ and D collapses naturally onto D'. Blue arrows denote a collapse.

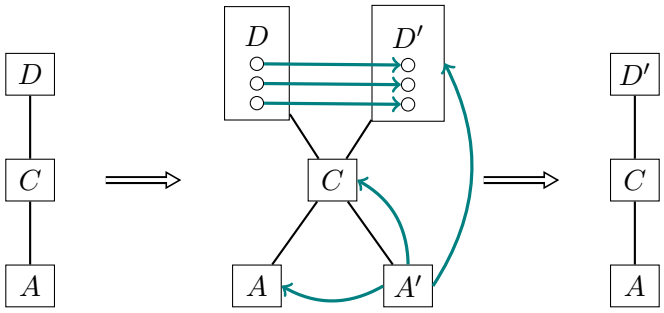

Lemma 8.17. *Let $\mathfrak{C}_{12}$ be constructed from some bipartite H_0 by one doubling and one collapse, as above. If $B = \emptyset$ and $E(A, D) = \emptyset$, then H_0 is collapsible onto $\mathfrak{C}_{12}$.*

Proof. Let the two steps in a construction of $\mathfrak{C}_{12}$ be such as in the statement. Recall that $G(S)$ denotes the induced graph on a vertex set S. Note that $H_0 = G(A \cup C \cup D \cup E \cup F)$ and that $\mathfrak{C}_{12} = G(A \cup C \cup D')$. For the following discussion cf. Figures 9 and 12.

Since $E(A, D) = \emptyset$, the graphs $G(A \cup C \cup D')$ and $G(A \cup C \cup D)$ are naturally isomorphic. Therefore, it is enough to show that it is possible to collapse $E \cup F$ onto $A \cup C \cup D$.

Let us write the collapse that produces $\mathfrak{C}_{12}$ as a homomorphism $f' : A' \cup D \cup E \cup F \cup F' \to A \cup C \cup D'$. By Lemma 8.16, either A' collapses naturally onto A or D collapses naturally onto D'.

Consider first that A' collapses naturally. We create a collapsing homomorphism $f : E \cup F \to A \cup C \cup D$ as follows:

- If $u \in E$ and $f'(u) \in A \cup C$, then $f(u) := f'(u)$. If $f'(u) = w' \in D'$, then $f(u) := w \in D$.
- For $u \in F$ with $u' \in F'$, if $f'(u') \in A \cup C$, then $f(u) := f'(u')$. If $f'(u') = w' \in D'$, then $f(u) := w \in D$.

We need to see that f is indeed a homomorphism, i.e., that all edges that touch $E \cup F$ are mapped onto edges of $G(A \cup C \cup D)$. To this end we make a case analysis:

- Since $G(E \cup F)$ is naturally isomorphic to $G(E \cup F')$ and $G(A \cup C \cup D)$ is naturally isomorphic to $G(A \cup C \cup D')$, the edges from $E(E \cup F, E \cup F)$ are preserved by f.

- Since $G(A \cup C \cup D \cup E)$ is naturally isomorphic to $G(A \cup C \cup D' \cup E)$, the edges from $E(E, A \cup C \cup D)$ are also preserved by f.
- Let $u \in F$, $v \in A$, $u \sim v$. Then $u' \sim v' \implies f'(u') \sim v \implies f(u) \sim v$, where we used that A' collapses naturally.
- Let $u \in F$, $v \in C$, $u \sim v$. Then $u' \sim v \implies f'(u') \sim v \implies f(u) \sim v$.
- Finally, let $u \in F$, $v \in D$, $u \sim v$. Then $u' \sim v' \implies f'(u') \sim v' \implies f(u) \sim v$.

Second, assume that D collapses naturally onto D'. In that case we give a collapsing homomorphism $f : \mathrm{E} \cup F \to A \cup C \cup D'$ as follows: if $f'(u) \in A \cup C$, then $f(u) := f'(u)$. If $f'(u) = w' \in D'$, then $f(u) := w \in D$. To see that f is a collapsing homomorphism, consider:

- Since $G(A \cup C \cup D)$ is naturally isomorphic to $G(A \cup C \cup D')$, f preserves the edges from $E(E \cup F, A \cup C \cup E \cup F)$.
- If $u \in E \cup F$, $v \in D$, $u \sim v$ consider the subcases (in all of them we use that D collapses naturally):
  - If $f'(u) \in A$, then $A \ni f'(u) \sim f'(v) = v' \in D'$, implying $E(A, D') \neq \emptyset$, a contradiction.
  - If $f'(u) \in C$, then $C \ni f(u) = f'(u) \sim f'(v) = v' \implies f(u) \sim v$.
  - If $f'(u) \in D'$, then $f'(u) \sim f'(v) = v' \implies f(u) \sim v$.

□

8.4 Putting things together

Proof of Theorem 8.2. By Lemma 8.9, if $\mathfrak{C}_{12}$ is constructible, there exists a construction of it by one doubling and one collapse starting from some H_0 that is not collapsible onto $\mathfrak{C}_{12}$ in the first place. But this is impossible by Lemmas 8.12 and 8.17. □

Remark 8.18. Our analysis, except for the computer-assisted part, does not depend on the number of vertices in $\mathfrak{C}_n$. Further program runs confirmed that also $\mathfrak{C}_{14}$ and $\mathfrak{C}_{16}$ are not constructible. On the other hand, one can see that $\mathfrak{C}_8$ and $\mathfrak{C}_{10}$ are constructible. ◊

References

[AS74] Miklós Ajtai and Endre Szemerédi. Sets of lattice points that form no squares. *Studia Scientiarum Mathematicarum Hungarica*, 9:9–11, 1974.

[BFS14] Harry Buhrman, Serge Fehr, and Christian Schaffner. On the parallel repetition of multi-player games: The no-signaling case. In *9th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2014)*, pages 24–35, 2014.

[BGKW88] Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: How to remove intractability assumptions. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, pages 113–131, 1988.

[Bod07] Hans L. Bodlaender. Treewidth: Structure and algorithms. In *Structural Information and Communication Complexity: 14th International Colloquium, SIROCCO 2007, Castiglioncello, Italy, June 5-8, 2007. Proceedings*, pages 11–25, 2007.

[BRR+09] Boaz Barak, Anup Rao, Ran Raz, Ricky Rosen, and Ronen Shaltiel. Strong parallel repetition theorem for free projection games. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques: 12th International Workshop, APPROX 2009, and 13th International Workshop, RANDOM 2009, Berkeley, CA, USA, August 21-23, 2009. Proceedings*, pages 352–365, 2009.

[BVY15] Mohammad Bavarian, Thomas Vidick, and Henry Yuen. Anchoring games for parallel repetition. arXiv:1509.07466, 2015.

[CCL92] Jin-yi Cai, Anne Condon, and Richard J. Lipton. On games of incomplete information. *Theoretical Computer Science*, 103(1):25–38, 1992.

[CWY15] Kai-Min Chung, Xiaodi Wu, and Henry Yuen. Parallel repetition for entangled k-player games via fast quantum search. In *30th Conference on Computational Complexity (CCC 2015)*, pages 512–536, 2015.

[Erd64] Paul Erdős. On extremal problems of graphs and generalized graphs. *Israel Journal of Mathematics*, 2(3):183–190, 1964.

[Fei91] Uriel Feige. On the success probability of the two provers in one-round proof systems. In *Structure in Complexity Theory Conference, 1991., Proceedings of the Sixth Annual*, pages 116–123, Jun 1991.

[Fei95] Uriel Feige. Error reduction by parallel repetition — the state of the art. Technical Report CS95-32, Weizmann Institute, 1995.

[FK91] Harry Furstenberg and Yitzhak Katznelson. A density version of the Hales-Jewett theorem. *Journal d'Analyse Mathématique*, 57(1):64–119, 1991.

[FRS88] Lance Fortnow, John Rompel, and Michael Sipser. On the power of multi-prover interactive protocols. In *Third Annual Structure in Complexity Theory Conference, 1988. Proceedings*, pages 156–161, 1988.

[FRS90] Lance Fortnow, John Rompel, and Michael Sipser. Errata for on the power of multi-prover interactive protocols. In *Fifth Annual Structure in Complexity Theory Conference, 1990. Proceedings*, pages 318–319, 1990.

[FV02] Uriel Feige and Oleg Verbitsky. Error reduction by parallel repetition — a negative result. *Combinatorica*, 22(4):461–478, 2002.

[HHC99] Chin-Wen Ho, Sun-Yuan Hsieh, and Gen-Huey Chen. Parallel decomposition of generalized series-parallel graphs. *Journal of Information Science and Engineering*, 15(3):407–417, 1999.

[HHM15] Jan Hązła, Thomas Holenstein, and Elchanan Mossel. Lower bounds on same-set inner product in correlated spaces. arXiv:1509.06191, 2015.

[HJ63] Alfred W. Hales and Robert I. Jewett. Regularity and positional games. *Transactions of the American Mathematical Society*, 106(2):222–229, 1963.

[Hol09] Thomas Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory of Computing*, 5(8):141–172, 2009.

[Pel95] David Peleg. On the maximum density of 0–1 matrices with no forbidden rectangles. *Discrete Mathematics*, 140(1–3):269–274, 1995.

[Pol10] D. H. J. Polymath. Density Hales-Jewett and Moser numbers. In *An Irregular Mind: Szemerédi is 70*, volume 21 of *Bolyai Society Mathematical Studies*, pages 689–753. Springer Berlin Heidelberg, 2010.

[Pol12] D. H. J. Polymath. A new proof of the density Hales-Jewett theorem. *Annals of Mathematics*, 175(3):1283–1327, 2012.

[Rao11] Anup Rao. Parallel repetition in projection games and a concentration bound. *SIAM Journal on Computing*, 40(6):1871–1891, 2011.

[Raz98] Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, 1998.

[Raz10] Ran Raz. Parallel repetition of two prover games (invited survey). In *Proceedings of the 25th Annual IEEE Conference on Computational Complexity, CCC 2010, Cambridge, Massachusetts, June 9-12, 2010*, pages 3–6, 2010.

[Raz11] Ran Raz. A counterexample to strong parallel repetition. *SIAM Journal on Computing*, 40(3):771–777, 2011.

[She88] Saharon Shelah. Primitive recursive bounds for van der Waerden numbers. *Journal of the American Mathematical Society*, 1(3):683–697, 1988.

[Sze75] Endre Szemerédi. On sets of integers containing no k elements in arithmetic progression. *Acta Arithmetica*, 27(1):199–245, 1975.

[Ver95] Oleg Verbitsky. The parallel repetition conjecture for trees is true. *Electronic Colloquium on Computational Complexity (ECCC)*, 2(13), 1995.

[Ver96] Oleg Verbitsky. Towards the parallel repetition conjecture. *Theoretical Computer Science*, 157(2):277–282, 1996.

[Wei13] Felix Weissenberger. Two-prover games for parallel repetition. Master's thesis, ETH Zurich, 2013.

A Listings of Computer-Assisted Proofs

Here we provide program codes for the computer-assisted proofs from Section 8. The programs are written in C++.

Listing 1: `construction.hpp` — The header file used by all programs.

```
#include <cassert>
#include <climits>
#include <cstdio>
#include <cstdlib>
#include <algorithm>
#include <vector>
using namespace std;

// Bitshifts have higher priority than comparisons.
// Comparisons have higher priority than bit operations.

// Mathematical modulo.
// Precondition: MOD > 0
inline int mod(int x, int MOD) {
  x %= MOD;
  return x + (x < 0 ? MOD : 0);
}

// Number of bits set to one in u.
struct PopCounter {
  int pcnt[1<<16];
  PopCounter() {
    assert(CHAR_BIT == 8 && sizeof(unsigned) == 4);
    for (int i = 1; i < 1<<16; ++i)
      pcnt[i] = pcnt[i/2] + i%2;
  }
} P;
inline int popcount (unsigned u) {
  return P.pcnt[u & ((1<<16)-1)] + P.pcnt[u >> 16];
}

// Undirected graph with set of vertices S.
// Invariant: all edges inside S.
struct Graph {
  unsigned S;
```

```
36    vector<unsigned> M;
37
38    Graph(const unsigned a_S,
39        const vector<unsigned>& a_M):
40      S(a_S), M(a_M) { }
41  };
42
43  // Doubles subset S of V(G).
44  // between[uprim] & (1<<u) indicates edge between
45  // u' in V(G') and u in V(G).
46  // Precondition: T is a subset of G.S
47  void double_graph(const unsigned T, const Graph& G,
48      Graph& Gprim, vector<unsigned>& between) {
49    const vector<unsigned>& M = G.M;
50    vector<unsigned>& Mprim = Gprim.M;
51    const int N = M.size();
52
53    Mprim.resize(N);
54    between.resize(N);
55    Gprim.S = T;
56
57    for (int u = 0; u < N; ++u)
58      if (1<<u & T) {
59        Mprim[u] = M[u] & T;
60        between[u] = M[u] & ~T;
61      } else Mprim[u] = between[u] = 0;
62  }
63
64  // Exchange vertices in T between G and G'.
65  // Precondition: T is a subset of G.S \cap Gprim.S
66  void exchange(const unsigned T, Graph& G, Graph& Gprim,
67      vector<unsigned>& between) {
68    vector<unsigned>& M = G.M;
69    vector<unsigned>& Mprim = Gprim.M;
70    const int N = M.size();
71
72    for (int u = 0; u < N; ++u) if (1<<u & T) {
73      unsigned old_M = M[u], old_Mprim = Mprim[u],
74        old_between = between[u];
75
76      M[u] = old_between & ~(1<<u);
```

```
77      for (int v = 0; v < N; ++v) {
78        M[v] &= ~(1<<u);
79        if (M[u] & 1<<v) M[v] |= 1<<u;
80      }
81
82      // It is important that 'between' has not been
83      // modified yet.
84      Mprim[u] = 0;
85      for (int vprim = 0; vprim < N; ++vprim)
86        if (u != vprim) {
87          Mprim[vprim] &= ~(1<<u);
88          if (between[vprim] & 1<<u) {
89            Mprim[u] |= 1<<vprim;
90            Mprim[vprim] |= 1<<u;
91          }
92        }
93
94      between[u] = old_M;
95      if (old_between & 1<<u) between[u] |= 1<<u;
96      for (int vprim = 0; vprim < N; ++vprim)
97        if (u != vprim) {
98          between[vprim] &= ~(1<<u);
99          if (old_Mprim & 1<<vprim)
100           between[vprim] |= 1<<u;
101       }
102   }
103 }
104
105 // Can Gprim be collapsed onto G?
106 // If yes, 'mapping' will contain a mapping
107 // from Gprim to G,  with mapping[uprim] == -1
108 // for uprim not in Gprim.S.
109 bool is_collapsible(const Graph& a_G,
110     const Graph& a_Gprim,
111     const vector<unsigned>& a_between,
112     vector<int>& a_mapping) {
113   struct RecursiveData {
114     const Graph& G;
115     const Graph& Gprim;
116     const vector<unsigned>& between;
117     vector<int>& mapping;
```

```
118     const vector<unsigned>& M;
119     const vector<unsigned>& Mprim;
120     const int N;
121
122     RecursiveData(const Graph& a_G,
123         const Graph& a_Gprim,
124         const vector<unsigned>& a_between,
125         vector<int>& a_mapping):
126           G(a_G), Gprim(a_Gprim), between(a_between),
127           mapping(a_mapping), M(G.M), Mprim(Gprim.M),
128           N(M.size()) {
129       mapping.resize(N);
130       fill_n(mapping.begin(), N, -1);
131     }
132
133     bool is_collapsible_rec(int uprim) {
134       if (uprim == N) return true;
135       if (1<<uprim & ~Gprim.S)
136         return is_collapsible_rec(uprim+1);
137       // invariant: u' < N and u' in V(G')
138
139       for (int u = 0; u < N; ++u) if (1<<u & G.S) {
140         if ((M[u] & between[uprim]) != between[uprim])
141           continue;
142         // invariant: u' -> u preserves edges between
143         // u' and G
144
145         bool ok = true;
146         for (int vprim = 0; vprim < uprim && ok;
147             ++vprim) {
148           if (Mprim[uprim] & 1<<vprim &&
149               !(M[u] & 1<<mapping[vprim])) {
150             ok = false;
151           }
152         }
153         if (!ok) continue;
154         // invariant: u' -> u preserves edges between
155         // u' and preceding vertices in G'
156
157         mapping[uprim] = u;
158         if (is_collapsible_rec(uprim+1)) return true;
```

```
159         }
160         return false;
161       }
162     } R(a_G, a_Gprim, a_between, a_mapping);
163
164     return R.is_collapsible_rec(0);
165   }
166
167   // Can G' be collapsed onto G such that both T and
168   // G'.S \setminus T do not collapse naturally?
169   // If yes, mapping will contain such mapping from
170   // G' to G, with mapping[u'] == -1 for
171   // u' not in G'.S.
172   // Precondition: T is a subset of G'.S which is
173   // a subset of G.S
174   bool is_unnaturally_collapsible(const unsigned a_T,
175       const Graph& a_G, const Graph& a_Gprim,
176       const vector<unsigned>& a_between,
177       vector<int>& a_mapping) {
178     struct RecursiveData {
179       const unsigned T;
180       const Graph& G;
181       const Graph& Gprim;
182       const vector<unsigned>& between;
183       vector<int>& mapping;
184       const vector<unsigned>& M;
185       const vector<unsigned>& Mprim;
186       const int N;
187
188       RecursiveData(const unsigned a_T, const Graph& a_G,
189           const Graph& a_Gprim,
190           const vector<unsigned>& a_between,
191           vector<int>& a_mapping):
192             T(a_T), G(a_G), Gprim(a_Gprim),
193             between(a_between), mapping(a_mapping),
194             M(G.M), Mprim(Gprim.M), N(M.size()) {
195         mapping.resize(N);
196         fill_n(mapping.begin(), N, -1);
197       }
198
199       bool is_collapsible_rec(int uprim) {
```

```
200         if (uprim == N) {
201           bool ok1 = false, ok2 = false;
202           for (int uprim = 0; uprim < N && (!ok1 || !ok2);
203               ++uprim) {
204             if (1<<uprim & ~Gprim.S) continue;
205             if (1<<uprim & T && mapping[uprim] != uprim)
206               ok1 = true;
207             else if (1<<uprim & ~T &&
208                 mapping[uprim] != uprim) {
209               ok2 = true;
210             }
211           }
212           return ok1 && ok2;
213         }
214
215         if (1<<uprim & ~Gprim.S)
216           return is_collapsible_rec(uprim+1);
217         // invariant: u' < N and u' in V(G')
218
219         for (int u = 0; u < N; ++u) if (1<<u & G.S) {
220           if ((M[u] & between[uprim]) != between[uprim])
221             continue;
222           // invariant: u' -> u preserves edges between
223           // u' and G
224
225           bool ok = true;
226           for (int vprim = 0; vprim < uprim && ok;
227               ++vprim) {
228             if (Mprim[uprim] & 1<<vprim &&
229                 !(M[u] & 1<<mapping[vprim])) {
230               ok = false;
231             }
232           }
233           if (!ok) continue;
234           // invariant: u' -> u preserves edges between
235           // u' and preceding vertices in G'.
236
237           mapping[uprim] = u;
238           if (is_collapsible_rec(uprim+1)) return true;
239         }
240         return false;
```

```
241       }
242     } R(a_T, a_G, a_Gprim, a_between, a_mapping);
243
244     return R.is_collapsible_rec(0);
245   }
246
247   // Precondition: T is a subset of G.S
248   inline unsigned neighbors (const unsigned T,
249       const Graph& G) {
250     const int N = G.M.size();
251     unsigned res = 0;
252     for (int u = 0; u < N; ++u) if (1<<u & T)
253       res |= G.M[u];
254     return res;
255   }
256
257   const int V = 12;
258   // Cycle with shortcuts C_V.
259   Graph original_G() {
260     Graph G((1<<V) - 1, vector<unsigned>(V));
261     for (int u = 0; u < V; ++u)
262       for (int s = -3; s <= 3; s += 2)
263         G.M[u] |= 1 << mod(u+s, V);
264     return G;
265   }
```

Listing 2: `non_empty_b.cpp` — Proof of Lemma 8.10.

```
1  #include "construction.hpp"
2
3  // Precondition: B, C disjoint, 0 in B
4  bool Bprim_filled(const unsigned a_C,
5      const unsigned a_B) {
6    struct RecData {
7      const Graph G;
8      const vector<unsigned>& M;
9      const unsigned C;
10     const unsigned B;
11     unsigned Bprim;
12     const int pB;
13     vector<int> B_list, Bprim_list;
```

```
14
15      RecData(const unsigned a_C, const unsigned a_B):
16          G(original_G()), M(G.M), C(a_C), B(a_B),
17          Bprim(0), pB(popcount(B)), B_list(pB),
18          Bprim_list(pB) {
19        for (int u = 0, ind = -1; u < V; ++u)
20          if (1<<u & B) {
21            ++ind;
22            B_list[ind] = u;
23          }
24      }
25
26      bool recursively_filled(int ind) {
27        if (ind == pB) {
28          // invariant: B, B', C (pairwise) disjoint
29          // invariant: edges of B and B' (inside and
30          // to C) isomorphic according to Bprim_list.
31          return is_rest_filled();
32        }
33
34        const int u = B_list[ind];
35        for (int uprim = 0; uprim < V; ++uprim) {
36          if (1<<uprim & (B|C|Bprim)) continue;
37          // invariant: uprim is "fresh"
38          if (M[uprim] & B) continue;
39          // invariant: no edges to B
40          if ((M[u]&C) != (M[uprim]&C)) continue;
41          // invariant: edges to C the same
42          bool ok = true;
43          for (int j = 0; j < ind && ok; ++j) {
44            const int v = B_list[j],
45              vprim = Bprim_list[j];
46            // a hack: '!' is used to convert to bool
47            if (!(M[v]&(1<<u)) != !(M[vprim]&(1<<uprim)))
48              ok = false;
49          }
50          if (!ok) continue;
51          // invariant: edges inside B and B' (so far)
52          // isomorphic
53          Bprim |= 1<<uprim;
54          Bprim_list[ind] = uprim;
```

```
55          if (recursively_filled(ind+1)) return true;
56          Bprim &= ~(1<<uprim);
57        }
58        return false;
59      }
60
61      // preconditions: B, Bprim, C disjoint
62      // B_list, Bprim_list, pB correctly filled
63      // B and B' isomorphic wrt each other and C
64      bool is_rest_filled() {
65        static Graph Gout(0, vector<unsigned>(V));
66        static vector<unsigned> between(V);
67        static vector<unsigned>& Mout = Gout.M;
68        static vector<int> mapping(V);
69
70        for (unsigned A = 0; A < 1<<V; A += 2) {
71          if (A & (B|C|Bprim)) continue;
72          // invariant: A, B, B', C disjoint
73          if (neighbors(A, G) & (Bprim)) continue;
74          // invariant: no edges between A and B'
75
76          const unsigned Dprim = ((1<<V)-1) &
77            ~(A|B|Bprim|C);
78          if (neighbors(Dprim, G) & (A|B)) continue;
79          // invariant: no edges between D' and A \cup B
80
81          Gout.S = A|Dprim;
82          for (int u = 0; u < V; ++u)
83            if (1<<u & A) {
84              Mout[u] = M[u] & A;
85              between[u] = M[u] & C;
86              for (int ind = 0; ind < pB; ++ind) {
87                const int v = B_list[ind],
88                  vprim = Bprim_list[ind];
89                if (M[u] & 1<<v) between[u] |= 1<<vprim;
90              }
91            } else if (1<<u & Dprim) {
92              Mout[u] = M[u] & Dprim;
93              between[u] = M[u] & C;
94              for (int ind = 0; ind < pB; ++ind) {
95                const int v = B_list[ind],
```

```
96                  vprim = Bprim_list[ind];
97                if (M[u] & 1<<vprim) between[u] |= 1<<v;
98              }
99            } else Mout[u] = between[u] = 0;
100
101         if (is_collapsible(G, Gout, between, mapping)) {
102           printf("FAILURE\nA = ");
103           for (int u = 0; u < V; ++u) if (1<<u & A)
104             printf("%d ", u);
105           printf("\n(B,B') = ");
106           for (int ind = 0; ind < pB; ++ind)
107             printf("(%d, %d) ", B_list[ind],
108               Bprim_list[ind]);
109           printf("\nC = ");
110           for (int u = 0; u < V; ++u) if (1<<u & C)
111             printf("%d ", u);
112           printf("\nD' = ");
113           for (int u = 0; u < V; ++u) if (1<<u & Dprim)
114             printf("%d ", u);
115           printf("\nmapping = ");
116           for (int u = 0; u < V; ++u)
117             printf("(%d->%d) ", u, mapping[u]);
118           printf("\n");
119           exit(0);
120         }
121       }
122       return false;
123     }
124   } R(a_C, a_B);
125
126   return R.recursively_filled(0);
127 }
128
129 // Assume E(A, D) is empty.
130 // Try all partitions of C_12 into A, B, B', C, D'
131 // s.t. in the last doubling:
132 // A is doubled and then A is fixed and A' collapsed.
133 // (non-empty) B is doubled and fixed together with B'.
134 // C is fixed in both steps.
135 // D is doubled, with D collapsed and D' fixed.
136 // Objective: show that resulting A', D cannot be
```

```
// collapsed onto the rest.
int main() {
  printf("non-empty B, V = %d\n", V);
  // Assume w.l.o.g. that 0 is in B.
  for (unsigned C = 0; C < 1<<V; C += 2)
    for (unsigned B = 1; B < 1<<V; B += 2) {
      // invariant: 0 in B
      if (B&C || popcount(B)%2 == 1) continue;
      // invariant: B, C disjoint
      if (Bprim_filled(C, B)) {
        // this should be never executed
        printf("INTERNAL ERROR\n");
        exit(1);
      }
    }
  printf("SUCCESS\n");
}
```

Listing 3: `non_empty_ad.cpp` — Proof of Lemma 8.11.

```
#include "construction.hpp"

// Assume B is empty and |E(A, D)| = 1.
// Try all partitions of C_12 into A, C, D' s.t. in the
// last doubling:
// A is doubled and A' is later collapsed.
// C is not doubled.
// D is doubled and later collapsed and D' is kept.
// Then try all choices for the edge between A and D.
// Goal: Show that resulting A', D  cannot be collapsed
// onto A, C, D'.
int main() {
  printf("|E(A,D)| = 1, V = %d\n", V);
  Graph G = original_G();
  for (unsigned A = 0; A < 1<<V; ++A)
    for (unsigned C = 0; C < 1<<V; ++C) {
      if (A&C) continue;
      // invariant: A, C disjoint
      unsigned Dprim = ((1<<V)-1) & ~(A|C);
      if (neighbors(Dprim, G) & A) continue;
      // invariant: no edges between A and D'
```

```
22
23        Graph tmp_G = G, Gprim(0, vector<unsigned>());
24        vector<unsigned> between;
25        vector<int> mapping;
26
27        double_graph(A|Dprim, tmp_G, Gprim, between);
28        exchange(Dprim, tmp_G, Gprim, between);
29
30        for (int u = 0; u < V; ++u) if (1<<u & A)
31          for (int v = 0; v < V; ++v) if (1<<v & Dprim) {
32            if (u%2 == v%2) continue;
33            // invariant: u and v do not create odd cycle
34            between[u] |= 1<<v;
35            between[v] |= 1<<u;
36            if (is_collapsible(tmp_G, Gprim, between,
37                mapping)) {
38              printf("FAILURE\nA = ");
39              for (int w = 0; w < V; ++w) if (1<<w & A)
40                printf("%d ", w);
41              printf("\nC = ");
42              for (int w = 0; w < V; ++w) if (1<<w & C)
43                printf("%d ", w);
44              printf("\nDprim = ");
45              for (int w = 0; w < V; ++w)
46                if (1<<w & Dprim) printf("%d ", w);
47              printf("\nu = %d v = %d\nmapping = ", u, v);
48              for (int w = 0; w < (int)mapping.size();
49                  ++w) {
50                printf("(%d -> %d) ", w, mapping[w]);
51              }
52              printf("\n");
53              exit(0);
54            }
55            between[u] &= ~(1<<v);
56            between[v] &= ~(1<<u);
57          }
58      }
59    printf("SUCCESS\n");
60  }
```

Listing 4: `natural_collapse.cpp` — Proof of Lemma 8.15.

```
#include "construction.hpp"

// Assume E(A, D) is empty.
// Try partitioning vertices of C_12 into A, C, D' s.t.
// in the last doubling:
// A is doubled and A' is later collapsed.
// C is not doubled.
// D is doubled and later collapsed and D' is kept.
// Objective: Show that every time either A' or D must
// be naturally collapsed.
int main() {
  printf("Natural collapse lemma, V = %d\n", V);
  Graph G = original_G();
  for (unsigned A = 0; A < 1<<V; ++A)
    for (unsigned C = 0; C < 1<<V; ++C) {
      if (A&C) continue;
      // invariant: A, C disjoint
      unsigned D = ((1<<V)-1) & ~(A|C);
      if (neighbors(D, G) & A) continue;
      // invariant: no edges between A and D

      Graph tmp_G = G, Gprim(0, vector<unsigned>());
      vector<unsigned> between;
      vector<int> mapping;

      double_graph(A|D, tmp_G, Gprim, between);
      exchange(D, tmp_G, Gprim, between);
      if (is_unnaturally_collapsible(A, tmp_G, Gprim,
          between, mapping)) {
        printf("FAILURE\nA = ");
        for (int u = 0; u < V; ++u)
          if (1<<u & A) printf("%d ", u);
        printf("\nC = ");
        for (int u = 0; u < V; ++u)
          if (1<<u & C) printf("%d ", u);
        printf("\nD = ");
        for (int u = 0; u < V; ++u)
          if (1<<u & D) printf("%d ", u);
        printf("\nmapping = ");
        for (int u = 0; u < (int)mapping.size(); ++u)
```

```
41              printf("(%d -> %d) ", u, mapping[u]);
42            printf("\n");
43            exit(0);
44          }
45        }
46    printf("SUCCESS\n");
47  }
```